Journal of Competition Law & Economics, 2023, **19**, 397–426
https://doi.org/10.1093/joclec/nhad009
Advance access publication date 19 September 2023
Article

OXFORD

RECOMMENDER SYSTEMS AND SUPPLIER COMPETITION ON PLATFORMS*

Amelia Fletcher**, Peter L. Ormosi*** & Rahul Savani#

ABSTRACT

Digital platforms can offer a multiplicity of items in one place. This should, in principle, lower end-users' search costs and improve their decision-making, and thus enhance competition between suppliers using the platform. But end-users struggle with large choice sets. Recommender systems (RSs) can help by predicting end-users' preferences and suggesting relevant products. However, this process of prediction can generate systemic biases in the recommendations made, including popularity bias, incumbency bias, homogeneity bias, and conformity bias. The nature and extent of these biases will depend on the choice of RS model design, the data feeding into the RS model, and feedback loops between these two elements. We discuss how these systemic biases might be expected to worsen end-user choices and harm competition between suppliers. They can increase concentration, barriers to entry and expansion, market segmentation, and prices while reducing variety and innovation. This can happen even when a platform's interests are broadly aligned with those of end-users, and the situation may be worsened where these incentives diverge. We outline these important effects at a high level, with the objective to highlight the competition issues arising, including policy implications, and to motivate future research.

KEYWORDS: Digital Platforms, Recommender Systems, Algorithmic Bias, Entry Barriers, Trustworthy Autonomous Systems

* The views expressed in this paper are the sole responsibility of the authors and cannot be attributed to the CMA, or Compass Lexecon, or any other parties.

** Non-Executive Director of the UK Competition and Markets Authority (CMA). Professor of Competition Policy, Norwich Business School, and Centre for Competition Policy, University of East Anglia, Norwich NR4 7TJ, UK. E-mail: amelia.fletcher@uea.ac.uk

*** Vice President, Compass Lexecon. Professor of Competition Economics, Norwich Business School, and Centre for Competition Policy, University of East Anglia, Norwich NR4 7TJ, UK. E-mail: p.ormosi@uea.ac.uk

Professor of Computer Science, University of Liverpool and The Alan Turing Institute, Department of Computer Science, University of Liverpool, L69 3BX, UK. E-mail: rahul.savani@liverpool.ac.uk

Received: November 29, 2022. **Revised:** June 6, 2023. **Accepted:** August 16, 2023

I. INTRODUCTION

Recommender systems (RSs) have become a key component of our online experience. These systems are responsible for choosing the products we see on our online retail journey; the results we get when we type in a particular search term; the songs that pop up on our favourite music streaming service; or the jobs that are advertised to us on our professional networking platform. They can vary widely both in form and in content. Some respond to specific search requests (for example, the consumer or end-user[1] tells an online booking platform that they want a hotel in a particular city); others inform the end-user without a search query (for example, a retail website listing products on their landing page). Some involve ranking of recommendations (for example, a retail platform displaying a list of products); some simply deliver what they have chosen for the end-user (for example, playlist continuation on music platforms or social media timelines). Some are truly personal, in that they depend on extensive information about the end-user, whereas others are more contextual and could potentially even be offered anonymously (for example, "you put this in your basket, so you might be interested in this").

One of the reasons that RSs have become so ubiquitous is that online end-users have access to a huge range of products and services on the platforms they use. In principle, end-users' search costs should be reduced by having easy access to so many products in one place, and this should lead them to make better decisions and enhance competition between suppliers (Goldfarb and Tucker, 2019). In practice, end-users struggle to choose from such a wide range of options, and use mental shortcuts, such as choosing the most prominent or highest-ranked option. RSs can help by providing relevant suggestions that reflect the end-user's preferences. Consumers have a strong propensity to click on such recommendations and to follow them in making their consumption choices.

Through this influence on end-user decision-making, RSs can also affect competition between the suppliers that use platforms to access those end-users. Of course, this impact can be positive. By enhancing the ability of end-users to choose across a wide range of offerings, and facilitating competition between suppliers, RSs can foster more efficient outcomes, increase allocative efficiency, and increase end-user welfare. There is a sizeable body of evidence on these benefits (Zhang, 2018; Waldfogel, 2017; Zentner *et al.*, 2013; Brynjolfsson *et al.*, 2011).

But what if RSs generate recommendations that do not fully reflect end-user preferences (assuming that these preferences are well defined)? The corollary to the positive effects of RS described above is that imperfections in the recommendations can potentially harm competition between suppliers. In general terms, this is obvious. If end-users buy products that are not especially suitable, or offer poor value for money, then this will be inefficient and weaken suppliers' incentives to offer end-users the best possible deals. Our focus in this article, though, is more specific. It is motivated by the finding from extensive computer science literature that RSs frequently generate recommendations that are not simply imperfect, but systemically biased. That is, they are biased in ways that are inherent to their design.

RSs typically do not have complete information about users' preferences and have to try and predict these on the basis of the data available to it. This process of prediction can, in turn, lead to systemic biases in the recommendations made. Of particular concern is the fact that RSs may favour particular types of suppliers or products over others and that this may, in turn, distort competition between suppliers and even—depending on the precise biases observed—drive increased concentration, raise barriers to entry and expansion, create market segmentation, and reduce variety. The impact on supplier incentives may lead to higher prices and reduced innovation.

[1] Hereafter, we use the term 'end-user' rather than 'consumer', reflecting usage of that term in the EU Digital Markets Act (Regulation (EU) 2022/1925).

This article is the first to review these issues in an overarching way. We take a high-level approach, considering the likely impact of such systemic biases from first principles. We do not include any modelling in this paper; albeit, we make reference to the limited literature available elsewhere. However, our core aim is to highlight the competition issues arising, including policy implications, and to motivate future economic and legal research.

Because our focus is on systemic biases, rather than strategic self-preferencing (which has been received greater focus in the economic literature to date), we assume broad alignment between the platform's interests and those of end-users, by focusing on RSs that seek to be 'customer-centric', in terms of delivering the most relevant recommendations that they can (we discuss the implications of relaxing this assumption in Section VII). For simplicity, and to isolate the impact of RS biases, we assume that consumers have well-defined and consistent preferences over the available options, but that they do not—absent the aid of an RS—know enough about these options to make a suitable choice.

The computer science literature highlights that, even where RSs are well intentioned in this way, they nonetheless tend to exhibit systemic biases. The precise biases observed will depend on the RS model design adopted ('Bias in Algorithm') and the data that feed into the model ('Bias in Data') (Abdollah and Mansoury, 2020). Data feedback loops can also be critical. The recommendations made by RSs are heavily reliant on the data they have available to them. If those data—or the extent of those data—are themselves biased, then this will tend to further bias these recommendations. But, of course, the data are typically derived from observing end-user reactions to previous recommendations made by the RSs. This creates a feedback loop whereby RS biases can become amplified over time by affecting the very nature of the data on which the RS is then trained.

The impact of RS biases may also be exacerbated by the 'choice architecture' through which recommendations are presented. Cognitive biases on the part of end-users can lead to their choices being unduly influenced by factors such as positioning or framing of options and their own selective interaction with the platform.

Although this article identifies and outlines these various potential effects of RS on supplier competition at a high level, we do not pretend to have a full understanding of their likely impact. Indeed, a key aim of the article is to motivate future research. We should note that we also do not consider the competition implications of RSs for competition between platforms themselves, which might, for example, arise due to differential access to data. This is an important topic, but beyond the scope of this article.

After examining the existing economic literature in this area (Section II), we provide initial background on the nature of RS (Section III) and describe the systemic biases inherent in RS design and implementation (Section IV). We then consider their likely implications for competition (Section V) and examine some ways in which platforms may seek to mitigate systemic biases (Section VI).

These sections all assume that platforms and end-user interests are broadly aligned and thus that RS are intended to be as customer-centric as possible. We then consider how a platform's commercial interests may, in fact, diverge from those of end-users and how this might be expected to affect RS design and thereby affect supplier competition (Section VII). Implications for competition policy (Section VIII) and wider policy (Section IX) are then considered before concluding.

II. EXISTING ECONOMICS LITERATURE

The computer science literature on RS is rich and documents a growing body of evidence on the market and social biases caused by RS. This literature is referenced where relevant throughout this article. However, far less has been written to date from an economics perspective about the implications of these biases for competition and market outcomes, especially in the context where platforms' interests are broadly aligned with those of end-users.

Most of the existing economics literature in this area focuses on the situation where platforms are incentivized to discriminate between different suppliers. Some of these examine 'self-preferencing', which can occur where a platform owns a vertically integrated supplier. Cure *et al.* (2022), Lee and Musolff (2021), and Kotapati *et al.* (2020) provide empirical analyses of this situation. On the theory side, Aridor and Gonçalves (2022) study the conditions under which a self-preferencing platform is welfare improving. Padilla *et al.* (2022) look at the conditions that make self-preferencing more likely, and De Corniere and Taylor (2014) show how integrated platforms (search engines) might be expected to distort rankings (search results). In addition, there are papers that consider platforms' profit incentives in relation to RSs absent vertical integration. Bourreau and Gaudin (2021) and Hunold *et al.* (2020) show how platforms may favour suppliers that are associated with cheaper content or higher revenue for the platform. Peitz and Sobolev (2022) find that platforms may deliberately recommend 'bad matches' to change the mix of consumers facing each supplier with a view to increasing surplus extraction. However, these models generally assume that platforms have perfect knowledge of consumer preferences and investigate only how this information is strategically used. The inherent complexities of real RSs, utilizing incomplete and imperfect data, are not considered.

There is also an economics literature on personalization more broadly. For personalized pricing, the literature largely pre-dates the widespread use of RS, and finds that personalized pricing may or may not be welfare enhancing, depending on conditions like market structure, whether firms compete on price or other product characteristics, the level of information asymmetry between end-users and suppliers, and the extent to which end-users exhibit cognitive biases (Ennis and Lam, 2021). There is less in the economics literature on personalization of attributes other than price, although Belleflamme and Peitz (2020) discuss the role and impact of online ratings, reviews, and recommendations on market competition. As discussed in Tucker (2012), there is a growing literature on the interaction between privacy and the use of RS, given the key role of personal data in personalizing recommendations. Hoffmann *et al.* (2013) discuss conditions under which the use of data for selective personalized disclosure tends to be good or bad for end-user welfare. Individualized recommendations or advertising can potentially also have an impact outside the RS. A field experiment by Fong (2017) finds that individually targeted advertising reduces search activity in other, nonadvertised products. Similar results are reported by Fong *et al.* (2019).

Most relevant for our work are a handful of papers that look specifically at how RSs affect specific market outcomes for suppliers selling through RSs, such as on sales volume or product diversity. As one might expect, RSs have been found to increase sales volume (Hosanagar *et al.*, 2014; Lee and Hosanagar, 2021). However, they can also affect the diversity of what is sold. They can increase the homogeneity of an individual end-user's consumption, a result that we refer to below as 'homogeneity bias' (Hosanagar *et al.*, 2014; Anderson *et al.*, 2020; Calvano *et al.*, 2023).

Greater homogeneity in individual consumption need not imply that diversity is reduced in aggregate. Holtz *et al.* (2020) show, through a field experiment on Spotify, that the RS increased the average number of podcast streams per user and decreased the average individual-level diversity of podcast streams, but increased the aggregate diversity of podcast streams. On the other hand, Fleder and Hosanagar (2009) find that RSs can cause overall diversity to fall, even if individuals themselves consume a more heterogeneous product mix.

Our article complements this previous literature but draws attention to a less researched area. Our focus is on the implications of systemically biased recommendations for competition between the suppliers that rely on a platform. To our knowledge, Calvano *et al.* (2023), Castellini *et al.* (2023), and Fletcher *et al.* (2023) are the only papers to date that directly investigate the impact of systemic RS biases on market outcomes. The former two papers simulate simple RSs and identify a tendency of the RS to increase market concentration and homogeneity. Calvano *et al.* (2023) focus on a single type of RS (collaborative filtering using matrix factorization), whereas Castellini *et al.* (2023) simulate and compare market outcomes under a wider set of RS types. The latter paper also demonstrates that RSs can increase entry barriers and compares the impact of RSs both for rational end-users and for those with limited attention. Fletcher *et al.* (2023) examine the pricing and quality decisions of suppliers under two different types of RS within a theoretical (Hotelling) framework.

Finally, we note that there is also a literature on the implications of RSs for online harms, for example, by fostering addiction, by creating echo chambers that can nurture conspiracy theories, or by defrauding end-users. These issues are important, but lie outside the scope of this article. We do, however, discuss how the presentation of RS results, through the chosen online choice architecture, can alter end-user choices and so affect competition.

III. RSs—WHAT ARE THEY?

This section is intended to provide a brief overview of the rudiments of RS design for readers without any background in RS. This should be sufficient to appreciate the points we make later about the impact of RS biases on supplier competition. Below, for notational simplicity, we will use the following terms: a *platform* is a digital arena that enables *end-users* to access *suppliers* for their *items*. The platform designs and operates the RS.

RSs are designed to provide recommendations for individual end-users. In a marketplace, these recommendations might be for goods; on a general search site, they might be for websites; and on a social media page, they might be for posts. For much of this article, we will assume that consumer preferences are well defined and consistent and that the platform designs RSs to be customer-centric, in the sense that RSs are designed with the intention to provide the most relevant recommendations for each individual end-user based on what the RS knows or predicts about that individual∼s preferences. Moreover, in our explanation, we implicitly assume that the end-users maximize their own utility by choosing which items to interact with or to not interact with any of the recommended items.[2]

In carrying out this function, RSs have to address a basic problem, which is that they do not typically start with complete information about end-users∼ preferences. Indeed, they may not even have full information on the characteristics of the items on offer. To address this problem, RSs essentially operate by collecting as much relevant data as possible∼often drawn from observing how end-users interact with items on the platform∼and using these data to ∼train∼ a statistical model. This model is designed to predict the items that are most likely to meet each end-user∼s preferences. More formally, RSs typically use known past interactions (for example, end-user A purchasing item x) to infer unknown future interactions (for example, whether end-user B will want to buy item y). Interactions between an end-user and an item can include any or all of clicks, likes, shares, time spent engaging, ratings, reviews, or purchases.

Importantly, the RS process typically involves a continuous feedback loop, as shown in Figure 1 below. First, data are collected. These data can relate to the end-users using the platform, the items on the platform, and interactions between those end-users and items. Second, these data are used to train, enhance, and utilize the RS model. Third, the recommendations generated by the RS model are fed back to the end-user. The way in which these are presented to end-users can be termed the ∼Online Choice Architecture∼, and how this is designed will itself

2 This is indeed in line with the assumptions made in Castellini *et*∼*al.* (2023).

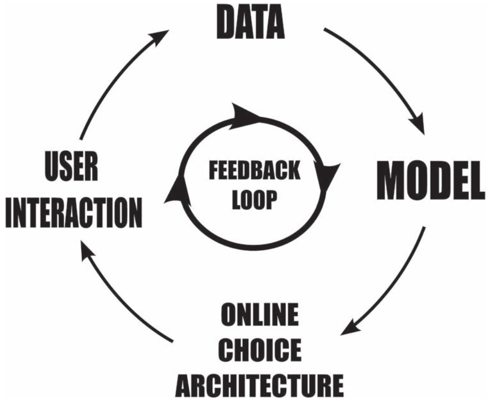

Figure 1. Recommender system pipeline.

tend to influence how the end-user reacts to the recommendations. Finally, the end-user~s consequential activity, based on the recommendations received, is fed back into the model as new data.

RSs do not necessarily make explicit recommendations. Often, they simply rank options, or place options in prominent positions, without specifically claiming they are recommendations. Sometimes, they will even make choices on behalf of end-users, such as a music streaming platform choosing what song to stream next. Either way, the nature of end-user behaviour is such that the end-users may be expected to behave ~as if~ the RS is making recommendations, and to choose from amongst those recommended options, so long as this gives them positive net utility.

Some RSs are entirely nonpersonalized, in that they do not use data on the end-user or their interaction with the platform. The most common nonpersonalized systems are popularity-based RSs, which simply recommend the most popular items (for example, news websites often use these methods, although some do also employ personalization). The remainder of this article focuses on personalized RSs, but it should be noted that even personalized RSs may be required to recommend options in the absence of relevant information about end-user preferences and, in such circumstances, may effectively revert to offering the most popular options.

RSs also vary in their complexity. Different RS designs may be optimal depending both on the quality and magnitude of data available and the nature of the recommendation required. For example, it may be more straightforward to provide an accurate recommendation for a relatively homogeneous product, such an AAA battery, where any differentiation is essentially vertical and end-user preferences are broadly aligned. It may be harder to provide recommendations for horizontally differentiated products, such as a shirt or a film, where end-users~ preferences are idiosyncratic and may diverge substantially.

A. Objectives of RS

A central element in RS design is the objective function that the RS is instructed to optimize. RSs can have differing objectives or even multiple objectives that are optimized jointly. Where RSs are intended to be customer-centric, the broad aim is to recommend items that give most value to end-users, but the more formal objective might be to optimize the relevance of the recommendation or to minimize the prediction error. This is not a trivial concept, however. To do this, most of the literature assumes that end-users~ preferences are well defined. This assumption allows one to measure the performance of RSs by looking at how far a prediction falls from the end-user∼s preference ordering. For the conceptual arguments in our paper, it is convenient to rely on this assumption. However we acknowledge that in the real-life setting of RSs, it is possible that end-users themselves are not fully aware of their own preferences or the ordering thereof or indeed that their preferences may be influenced by the RS.

Regarding the objective function of the RS, it can be the case that there is a tension between offering the most relevant recommendations in the short term and enhancing the quality of recommendations over the longer term. To fully understand each end-user∼s preferences, RSs may need to make some less obvious recommendations, as end-users∼ reactions to these can generate more information on their preferences than their reactions to the more obvious recommendations. As a simple example, a music streaming site is unlikely to ascertain that someone has a penchant for jazz if it only ever recommends pop. This means that it might be in an end-user∼s own long-term best interest that they sometimes receive apparently less relevant recommendations in the short term. Such experimentation can also reduce the biases inherent within RSs.For some types of products, end-users may gain positive benefit from an element of novelty or diversity in the recommendations provided (Zheng and Wang, 2022). Moreover, an RS can itself form and reshape end-user preferences. An end-user with a set of preferences at a given time might have different preferences the next day, and the change may be at least partially due to the RS.

Finally, it should, of course, be noted that platforms are typically commercial entities, seeking to maximize profit. While platforms∼ incentives may be well aligned with those of their end-users because they make more profit by providing more relevant recommendations∼this need not always be the case. Platforms may incorporate commercial objectives such as revenue or profit maximization.

We discuss the implications of such commercial objectives in Section VII.

IV. RS AND SYSTEMIC BIASES IN RECOMMENDATIONS

Our primary focus in this article is the situation where the interests of a platform and its end-users are broadly aligned, in the sense that the platform seeks to provide customer-centric recommendations. Even where this is the case, the design and implementation of RSs is invariably imperfect in practice. To some extent, this is unsurprising; RSs do not have full information on individual preferences, and these are inherently hard to predict.

What is perhaps less obvious is the fact that, in seeking to address this issue of incomplete information, RSs can exhibit systemic biases. We use the word bias here as a shorthand expression to refer to recommendations that deviate from the recommendation that would maximize the user's utility. Examples of recommendation bias could include recommendations that disproportionately feature popular items as opposed to items that an end-user would value highest, or recommendations that feature items with more extensive metadata over those that would maximize an end-user's utility. Such biases can arise from both the design of RS models and from the data feeding into them. There is substantial computer science research in this area, usefully surveyed in Chen *et al.* (2020).

In this section, we draw on this body of literature to outline the key biases arising.[3]

A. RS Design and Systemic Biases in Recommendations

There are two canonical forms of model that are most frequently used in RS: 'collaborative filtering' and 'content-based filtering'. Beyond these, there are other methods, such as 'knowledge-based filtering'. These various methods can be used on their own, or in combination as a 'hybrid RS'.

[3] A comprehensive overview of the RS discussed below can be found in Aggarwal *et al.* (2016).

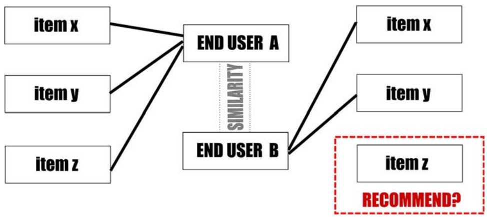

Figure 2. Collaborative filtering—a simple example.

(a) Collaborative filtering is one of the most widely used methods in RS. Collaborative filtering makes predictions about the likely interactions of a user with items that this user has not yet interacted with, based on the past interactions of this and many other users. This draws on the intuition that end-users who have similar preferences as regards one item will be more likely to have similar preferences more generally. For a given active user, the RS effectively then recommends items that those similar users have liked but which the active user has not yet interacted with. In other words, recommendations are effectively of the form '*users who have liked similar items to you in the past also like this item which you haven't tried yet*'.

Figure 2 provides a visual illustration. If end-users *A* and *B* both like items *x* and *y*, and end-user *A* also likes item *z*, then the RS infers that item *z* is a good recommendation for end-user *B*.

There are various forms of collaborative filtering. 'Neighbourhood techniques' identify similar users, or 'neighbours', in preference space. Alternatively, they can find neighbourhoods based on item similarity. In either case, these methods construct similarity measures between users (or items) and recommend content based on these measures. Item–item collaborative filtering can be preferable to use–user collaborative filtering when there are more users than items because, in this case, each item tends to have more interaction data than each user, so an item's average rating is more stable than a user's rating profile.

Another form of collaborative filtering involves 'matrix factorization'. The RS creates a matrix of users and items, which includes all known information from past user–item interactions. This matrix inherently includes many blank cells (items that given users never interacted with). The model effectively predicts values for these missing cells by assuming that this matrix can be factorized into two smaller matrices—one that represents a number of key attributes for each item and one that represents the preferences over those key attributes for each end-user—and estimating the missing cells on the basis of the cells that are complete.

Collaborative filtering has some appealing characteristics. For example, it does not need information on the characteristics of the end-users or of the items, and it can be applied irrespective of the domain, whether it is retail or music recommendation. On the negative side, however, it suffers from a number of systemic biases in the recommendations made.[4]

Most prominent of these is **popularity bias**. This arises because collaborative filtering works best for relatively highly purchased items and less well for those in the 'long tail' for which there is relatively sparse data. As a result, more popular items tend to be recommended more than is proportionate, given end-users' underlying preferences. Collaborative filtering also suffers from the 'cold start' problem, whereby items that have never previously been purchased rarely get recommended and so rarely get purchased. In economic terms, this is a form of **incumbency bias**. Established products are recommended simply because the RS has more data on them,

4 Biases in collaborate filtering methods are well known and widely documented in the computer science literature, see, for example, Yao and Huang (2017), Sun *et al.* (2019), Bobadilla *et al.* (2012), and Guo *et al.* (2014).

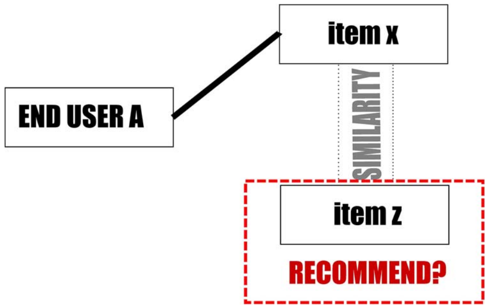

Figure 3. Content-based filtering—a simple example.

not because they are the best. New products are not recommended, even if they are better fit with end-users' underlying preferences.

(b) Content-based filtering. Content-based filtering analyses the features of items to identify similarities between them and then links this to data on the items the end-user has interacted with; in contrast to collaborative filtering, where item similarity is based only on past interactions between users and items, content-based filtering uses a rich set of features about the items themselves, such as the type of movie or leading actors, or a song's artist or style. For example, as shown in Figure 3, if the end-user *A* has previously bought item *x*, and items *x* and *y* have similar features, then the RS infers that end-user *A* will also like item *y*. An obvious (if over-simplistic) example might be where an end-user likes one song by Adele, and the RS therefore recommends another song by Adele.

An advantage of content-based filtering is that, unlike collaborative filtering methods, it is better suited to addressing the 'cold start' problem. Content-based filtering can be based on a diverse range of item features and can even use natural language processing (NLP) techniques to extract features from item descriptions. This enables it to be applied widely, even to brand-new products.

On the other hand, content-based filtering can also generate systemically biased recommendations.[5] First, it tends to create filter bubbles and make recommendations that are disproportionately similar to those already consumed. This **homogeneity bias** in recommendations at the individual level can fail to reflect the diversity that exists within an end-user's true underlying preferences or indeed any specific preference the end-user might have for variety.

A second limitation of content-based filtering is that it is weak in handling new end-users who have not yet purchased many items. For such end-users, and depending on the precise RS design, it may tend to recommend products it knows to be generally popular (**popularity bias**). While such recommendations may be right for some end-users, they will be systemically biased for others, relative to their underlying preferences.

(c) Other methods: *Knowledge-based filtering* is based on interactively specified user requirements, rather than data on user–item interactions. Such information can be requested from end-users—such as Pinterest asking new users for their interests—or acquired from outside the system. It can also include information derived from the end-user's activity on the wider platform. The data are used to build a 'profile' of the end-user, identifying a number of individual characteristics. An advantage of knowledge-based filtering is that it can utilize substantial data to develop its recommendations, and new or irregular end-users may be able to input (or port)

[5] Once again, these biases are widely documented. For example, see Nguyen *et al.* (2014), O'Callaghan *et al.* (2015), Nagulendra and Vassileva (2014), and Pagano *et al.* (2016).

sufficient data into the RS to generate reasonable recommendations. For the same reason, knowledge-based filtering methods are more suitable for nonroutine, infrequent purchase decisions (for example, house or car purchase).

However, knowledge-based filtering can also lead to systemically biased recommendations (Mandl *et al.*, 2011). The specified criteria are frequently not sufficient to narrow down to only one option, so a choice of recommendation then still needs to be made. This might be the most popular item that meets all the criteria, but if so, this recreates at least some tendency towards **popularity bias**. Moreover, if individuals with similar user requirements are always shown similar content, this can lead to disproportionate within-group conformity or **conformity bias**. In a content-based environment, this effect can create 'echo chambers' or 'filter bubbles'.

Another RS method is *demographic filtering*, where demographic information about the user is leveraged to learn classifiers that can map specific demographics to ratings or buying propensities. A relatively sophisticated form of this is *social filtering*, which utilizes information from social networks to make recommendations, sometimes of the form 'others in your network like this product' and sometimes even naming the individuals concerned (Yang *et al.*, 2014). Finally, *utility or propensity-based RSs* use the available data to estimate individuals' utilities of, or propensities towards, particular items. A key issue with these RS is that they tend to require very extensive data.

(d) Hybrid models can combine any of the above methods. For example, Zanker & Jessenitschnig (2009) describe how collaborative and knowledge-based filtering can usefully be combined, while De Campos *et al.* (2010) discuss the combination of content-based and collaborative filtering. Hybrid methods can also combine the above methods with simple non-personalised methods, such as popularity recommenders. For example, content-based filtering might be used to reduce the recommendation set, with the most popular of this set then being recommended. Methodologically this can be done through ensemble methods, which layer the output of multiple algorithms, or monolithic systems, whereby an integrated recommendation algorithm is created that incorporates different forms of filtering.

B. Data and Systemic Biases in Recommendations

It will be clear from the discussion so far that RSs are inherently data-hungry. The more data they have, the more accurately they can predict end-user preferences and therefore make relevant recommendations. The data used in RSs can be explicit (whereby users are asked to provide certain information about themselves or the items), or implicit (whereby the user's interaction with items on the platform is recorded).[6]

The heavy reliance by RSs on data means that biases in data inputs can add to the intrinsic biases within RSs discussed above. These data biases can take two main forms: the data can be accurate, but their extent (their sparsity) may be biased, or alternatively, the data themselves may be biased.

Sparsity bias. Even if the data available to RS are accurate, there may be strong variation in the density of data, with data sparser in certain areas. This is important because RSs of any sort are typically less likely to recommend items to end-users where they have only sparse data to support that recommendation. If there are particular categories of item, end-user characteristic, or item–end-user interaction for which data are especially sparse, this can lead to the RS generating biased recommendations. In particular, it can be this data sparsity effect that drives RS to favour both

[6] This need for extensive data adds to the already well-documented tendency of platform markets to become concentrated; the largest RS will have more data that enable them to make more accurate recommendations. Since end-users value more accurate recommendations, smaller platforms struggle to compete, even if their RS is theoretically better, without access to such extensive data. This situation risks being worsened if smaller platforms are targeted by weaker suppliers, who struggle to get their products recommended within more accurate RSs, because this will tend to make the smaller platforms still less attractive to end-users. Our focus in this article, however, is not on the implications of RS for the platform market itself.

popular products and established products, given that the RS is likely to have more extensive data on them, leading to the **popularity** and **incumbency biases** discussed above (Idrissi and Zellou, 2020).

Sparsity bias can arise naturally. For example, some items are intrinsically 'long tail' items, with limited appeal and thus are likely to generate limited interaction data. However, sparsity bias can also result from more behavioural considerations. First, end-users are disproportionately likely to review or click on items that they very much like or very much dislike (Marlin *et al.*, 2012; Schoenmueller *et al.*, 2020). They are less likely to bother rating 'more average' products (say, worthy of three stars out of five) or to click on content they might find interesting but is neither exciting or alarming. Second, certain types of end-users may be more likely to leave reviews than others. For example, more men review films on movie sites than women. Absent specific adjustments to address this fact, it might, in turn, be expected to skew the rankings of films on those sites towards films that men tend to prefer.[7] Third, in terms of item features, concerns have been expressed that NLP data collection methods work less well for non-English language content, and thus that data may be sparser for such content. Fourth, in some markets, smaller suppliers may struggle to provide all of the data required by the RS, meaning that data on their product will be partial.

The key implication of each of these types of sparsity bias is that the availability of data will be nonrandom. There will be disproportionately more data on products that elicit extreme reactions, or on films that men tend to like, or on English language songs, or on products for which the supplier is better positioned to provide the required data. If there are more data, then those products will be more likely to be recommended, leading to recommendations that are biased relative to end-users' true preferences.

Herding bias. So far, we have assumed that the data inputs used by RS are accurate, but this need not be true either. For example, it has been shown that users tend to exhibit 'herding bias' when rating products, whereby they may raise their ratings of a product if they see that others have rated it highly, or conversely lower them if they see that others have a negative opinion (Lederrey and West, 2018; Liu *et al.*, 2016; Wang and Wang, 2014). This can lead to data being collected that diverge from end-users' underlying preferences in a nonrandom way. Such herding effects can potentially exacerbate some of the systemic recommendation biases identified above, such as popularity bias, incumbency bias, homogeneity bias, and conformity bias. Herding can be especially problematic if the initial ratings or reviews are themselves biased (as might be the case if they have received paid-for 'fake reviews').

C. Feedback Loops

Data may also become biased through the effect of feedback loops. As discussed in Section III, the RS process works as a continuous loop, with data that are collected on the platform being fed back into the RS design and application. This is important because if there is bias at any stage of the loop, this will affect the data being generated and fed back into the process. For example, suppose an end-user is recommended the most popular product by an RS, not because this is the most suitable product for that end-user but because the RS exhibits some popularity bias. If the end-user accepts the recommendation, then this will both add to the extent of data available on this item and add to the average positivity of data about this item. This will be true even if there are a hundred other items that the end-user would have preferred, had they only known about them. But both of these factors will tend to increase the likelihood that the item is recommended again, thereby exacerbating the initial popularity bias.

7 Marcus Beard, "IMDB analysed: how do men and women's favourite films differ?, One Room With A View", 10 August 2016. https://oneroomwithaview.com/2016/08/10/ imdb-analysed-men-womens-favourite-films-differ.

Within such data feedback loops, even small initial recommendation biases can quickly become more extreme (Mansoury *et al.*, 2020; Jannach *et al.*, 2015). Data feedback loops can have an especially strong effect in driving homogeneity in recommendations (Chaney *et al.*, 2018). When the nature of the data feeding into RS has been influenced by the nature of the RS in this way, this is sometimes known as 'algorithmic confounding'. As a platform grows in users and in content, it may well be the case that ever less of the data available to the RS will be 'organic' (unaffected by the RS design). Another type of feedback loop can occur when the recommendations made by an RS actually act to shift users' preferences over time. This can potentially lead to even greater market homogeneity (Mansoury *et al.*, 2020).Of course, if the RS has sufficient upfront knowledge about the users and their preferences (for example, by having a near-complete matrix of user–item reviews), then the RS model will be less likely to change in response to the data that are created and thus be impact less by data feedback loops (albeit this can still change if new end-users or items are introduced into the system).

D. Choice Architecture

The feedback loops described above may also be influenced by the way in which recommendations are presented to end-users—the so-called 'choice architecture'. RSs are designed to generate an ordered list of relevant recommendations, but this list then needs to be presented to end-users for them to make their choice. This presentation can take a variety of forms. For example, end-users could be given a single recommendation or a list of 10 to choose from. If they are presented with more than one option, these may be randomly ranked or ranked in order of relevance. Some may be more prominent than others, or end-users may be given other forms of 'nudge' towards particular options.

The choice architecture adopted will tend to affect end-user choices, due to a variety of well-evidenced end-user behaviours, sometimes referred to as 'behavioural biases'. This is important not only because it affects the suitability of those end-user choices but also because it can, in turn, affect the data that are fed back into the RS. For example, **exposure effects** result from end-users only engaging and thus providing data on those items to which they are exposed. Given the huge number of items typically available on platforms, end-users tend to focus only on the small number of items that they are shown, unless they have a specific reason to dig deeper into the long tail. A similar concept is **positioning effects**, which result from the fact that end-user interactions are also influenced by rankings and prominence, with end-users far more likely to click, and thus generate data, on items at or near the top of rankings or prominent on the screen (O'Brien and Keane, 2006; Ursu, 2018; Ghose *et al.*, 2014; Carare, 2012; Aguiar and Waldfogel, 2021).

The way in which recommendations are framed can also be critical (**framing effects**). For example, end-users may be more inclined to purchase where recommendations are accompanied by particular phrases. These might include 'people like you bought…' (drawing on 'social proof'); 'only three items left…' (drawing on 'scarcity effects'), or '30 per cent off…' (drawing on 'reference point effects'). Likewise, Alexa 'Hunches', which involve Amazon's virtual assistant Alexa making private recommendations to people in their homes, plug into people's routines to then make recommendations in daily situations real time, without people even needing to engage with the idea of searching. This is likely to produce different outcomes from more 'passive' RSs where the end-user takes a more proactive role. Indeed, the choice architecture can itself be directly influenced by the RS (Jesse and Jannach, 2021). For example, movie streaming platforms can personalize the film artworks or thumbnail that each viewer sees on their TV (an end-user who likes horror films may be shown more gruesome thumbnails for each film).

While the impact of choice architectures on end-user choices may be unavoidable, given the difficulties end-users face in digesting huge amounts of information, they will nonetheless tend to influence the data collected and thus the data feedback loops described above. In general, a choice architecture that is more balanced, and better designed to encourage active choice amongst end-users between a set of alternatives, is likely to not only deliver better outcomes for end-users in the short term but also generate less biased data and thus less biased recommendations in the longer term.

By contrast, choice architectures that give end-users less choice—whether this is explicit or more implicit (for example, through the use of rankings)—will tend to exacerbate the systemic recommendation biases of RS models described in Section IV.A. Castellini *et al.* (2023) show the market consolidating impact of platforms changing the choice architecture, to give more prominence to certain items when users have limited attention and are more likely to engage with items ranked at the top of a recommendation list.

The discussion above has shown that recommendation biases can stem each of the different stages of the RS pipeline described in Figure 1 above, and indeed, it may be very difficult to unravel exactly what first caused any identified bias. Moreover, if an RS model is initially biased, it may prove very difficult to fully de-bias recommendations at later stages, as the RS model will by that point be learning from data collected when it was biased.

V. IMPLICATIONS OF SYSTEMIC RS BIASES FOR COMPETITION BETWEEN SUPPLIERS

Much of the above discussion on the systemic biases associated with RSs has been covered extensively in the computer science literature. However, that literature does not go on to consider the economic implications of systemic RS biases for competition between suppliers who are dependent on RSs for their sales.

In this section, we argue that systemic biases in the RS deployed on digital platforms will inevitably affect competition on the supplier side of the market. This is because there is a close link between recommendations and actual sales. Consumers have a strong propensity to click on recommendations (De los Santos and Koulayev, 2017; O'Brien and Keane, 2006; Joachims *et al.*, 2005) and to make consumption choices on this basis (Ursu, 2018; Ghose *et al.*, 2014; Carare, 2012; Aguiar and Waldfogel, 2021; Lee and Musolff, 2021). Thus, if an RS produces systemically biased recommendations, this will tend to drive demand towards those products that benefit from the bias and away from those that do not. This will, in turn, distort or dampen competition between the suppliers of those products.While the direct effects of any such distortion are short term, there may be a long-term impact, too. This is partly due to the data feedback loops discussed above. If biased RSs give rise to biased end-user choices, data that are then fed back into the RS, then these biases risk becoming more extreme over time. However, suppliers' long-term incentives are also relevant. Suppliers will have less incentive to invest or innovate into products, if those products are less likely to be recommended.

Of course, RS biases need not always be a problem. In some situations, they can, in fact, mitigate other market failures and thereby improve market outcomes. For example, popularity bias has been shown to be useful in ensuring that the highest-quality products come out on top (Ciampaglia *et al.* 2018; Zhao *et al.* 2021).

It is also important to consider the relevant 'counterfactual'. Since platforms inherently have incomplete information, there is unlikely to be any such thing as a fully unbiased RS. Moreover, we might well observe similar biases occurring in an offline world, albeit in a less direct way. The fact that physical retail stores can only stock a limited variety of products, for example, could generate a form of 'popularity bias'. As such, even where RS exhibit systemic biases in their recommendations, these may be less serious than would have been the case in an offline environment.

Nonetheless, in general, we would expect competitive outcomes to be enhanced through recommendations that more fully reflect underlying end-user preferences, and therefore competition to be harmed if recommendations deviate from this in a systemic way.

The harm to competition arising from such RS biases may be expected to be particularly serious where the platform concerned itself faces limited competition, for example, due to network effects and scale economies, or where it has 'single homing' end-users. These factors tend to make platforms an essential route to end-users, giving them strong 'bottleneck' market power over suppliers (Armstrong and Wright, 2004; Coyle *et al.*, 2019). If RS biases affect competition on such platforms, this is more likely to affect competition across the upstream supply market as a whole.

In the discussion below, we consider a situation where there is indeed a single platform, providing the only means of access for suppliers to end-users. This assumption is for simplicity. In general, the same effects are likely to arise with multiple platforms, so far as they have a degree of 'bottleneck' market power over access to end-users. However, the impact of each individual platform on overall competition between suppliers would tend to be reduced with multiple platforms.

It is important to emphasize that the concerns raised at this stage are not associated with platforms acting strategically to limit or distort competition. For the time being, we are still assuming that the incentives of platforms and end-users are broadly aligned, in the sense that RSs are intended to be as customer-centric as possible. Our focus is on the anticompetitive effects that may result from the inherent characteristics of RSs. The choice of RS design may seem like a technical question, but it contributes directly to the type and extent of any such impact on competition.

In this section, we will highlight a number of likely implications for supplier competition of the systemic recommendation biases we have identified in RSs. These have been developed at a high level, on the basis of first principles. However, we also highlight those specific implications that have been tested using simulation techniques or theoretical work.

A. Increased Market Concentration and Supplier Incentives

The first concern is that popularity bias can drive markets towards becoming more concentrated, potentially with just one or two very popular products in each category and no other products able to gain recommendations, and therefore end-users. This sort of 'blockbuster' effect is well known in certain sectors, such as the movie industry, but has hitherto been less of an issue in markets more generally. Moreover, the 'blockbuster effect' for movies is inherently short lived. When transposed into less dynamic markets, it risks generating serious long-term concentration in supplier markets. This impact of popularity bias is sometimes known as the 'Matthew' effect or the 'Rich get richer' effect.

The tendency towards market concentration has important implications for supplier incentives too (Calvano *et al.*, 2023; Castellini *et al.*, 2023; Fletcher *et al.*, 2023). If more popular firms have an advantage in terms of being recommended by the RS, then this may allow them to command a price premium relative to new entrant or less popular firms or to reduce their investment in quality. The situation is more complex for rival (challenger) suppliers. For those, a key question is how easy it is to contest the position of 'most popular' product through pricing low or investing in high quality and thereafter gain the benefits that such popularity brings. If this is possible, these rival suppliers may actually offer lower prices and higher quality to gain pole position than they would with a less biased RS (although average prices are still likely to be higher). By contrast, if it is almost impossible for a rival to gain the position of 'most popular' product, the impact of 'popularity bias' in the RS may dampen the competitive incentives for rivals too, with prices increasing for all suppliers.

The online choice architecture through which recommendations are presented to end-users can play a key role here, either exacerbating or mitigating the impact of the popularity bias within the RS. If there are a small number of recommendations (as in the Amazon 'Buy Box', for example) or if there is a clear ranking, then the top recommendation is likely to gain more attention, user interaction, or sales. This will, in turn, generate data that support the popularity of that product, creating a data feedback loop that increases the drive towards market concentration. By contrast, if the top four recommendations are given equal prominence on the screen (as in Google's Shopping box, for example), with the precise positions of the four randomized, then this may have a less concentrating effect on end-user choices and competition.

As discussed above, data-related 'sparsity bias' is also relevant. If end-users are disproportionately more likely to rate, review, click, or share the very best items, relative to those that are only slightly less good, this will tend to exacerbate popularity bias and its associated impact on competition.

B. Increased Barriers to Entry and Expansion

Another competition risk arising from RS systemic biases is that they may create barriers to entry and expansion. If popular or well-established items feature disproportionately in recommendations (due to 'popularity' and 'incumbency bias'), it will be more difficult for smaller or new entrants to be recommended, even if they are better or more innovative than the incumbents. Since recommendations are often critical for gaining end-users, this can constitute an important barrier to entry or expansion (Castellini *et al.,* 2023). Barriers to entry and expansion are of particular concern when they limit the potential for new innovative products (or versions of products) to gain end-users. This risks reducing valuable innovation and market dynamism, which have, to date, been a key positive aspect of the internet.Data feedback loops can further worsen these barriers, as the RS is more likely to collect data on interactions from popular and incumbent items as opposed to new entrants. This relative lack of data being collected on new or smaller entrants makes it increasingly difficult over time for them to climb the rankings within RS and thereby gain custom. As before, the online choice architecture around how the recommendations are presented can amplify the impact of RS data feedback loops.

Entry barriers are likely to be higher under RS models such as collaborative filtering, which normally rely on data about past user–item interactions, than under content-based or knowledge-based RS, which utilize data that can be collected separately.

C. Increased Homogeneity and Reduced Variety

As discussed above, RS can exhibit 'homogeneity bias', whereby people tend to be shown the same products, or the same sorts of products, over time. In essence, once an RS thinks it has learned an end-user's preferences, it will seek to meet those preferences, and the consequent lack of diversity in its recommendations means that it never learns more about that end-user's preferences.

This can be detrimental for end-users if they, in fact, have wider preferences than the RS model identifies or if they simply value variety for its own sake. However, it also has implications for competition.

These competition implications may well be negative. For example, if RSs exhibit 'popularity bias', this will tend to reduce sales of items in the long tail, which will reduce variety and increase homogeneity. If this leads to suppliers ceasing to provide those products, this will be harmful to those end-users who value this variety. This is more serious than it may sound because the growth of the long-tail sales has been one of the successes of the digital economy to date and a key driver of growth (Brynjolfsson *et al.,* 2011). Any deterioration of its viability would harm the economy and end-users alike. In general, end-users that most like 'long tail' products appear to be least well served by RS (Abdollahpouri *et al.*, 2019; Kowald *et al.*, 2020).Likewise, innovative new products will tend to be somewhat heterogeneous relative to existing products—this is indeed what is good about them. If 'homogeneity bias' in RSs reduces their ability to gain sales, then it is less likely that the required R&D will be funded in the first place, to the detriment of those end-users who would have benefited from these new products.

However, increased homogeneity can potentially have a positive impact on competition too. If end-users face a more homogeneous sets of choices, then suppliers will need to compete harder on price to win custom, leading to a reduction in equilibrium prices (Calvano *et al.*, 2023). Nonetheless, there may be downsides even to this apparently positive competition, since intensive price competition of this sort can act to reduce competition in other dimensions such as range, quality, service, or innovation.

D. Increased Market Segmentation

If specific categories of end-users are given unduly similar recommendations ('conformity bias'), this can lead to 'filter bubbles' or 'echo chambers'. While policy concerns in this area typically focus on the cultural implications, there are also potential competition risks. In particular, products can become identified with one social group and therefore never recommended to end-users in another social group.

On the positive side, such market segmentation can potentially act to increase market-wide variety, even if each individual group receives more homogeneous recommendations (Holtz *et al.*, 2020). However, it can also reduce competition within each social group, since, from a particular social group's perspective, item *x* may not viewed as a substitute for item *y*, even if they effectively fulfil the same function. Prices can potentially rise as a consequence of this 'within-social group concentration' (Fletcher *et al.*, 2023), although, in a somewhat different context, Gal-Or and Gal-Or (2005) noted a similar effect. At an extreme, individual social groups might form distinct relevant markets, with different suppliers and conditions of competition. The difficulty involved in having to break into a series of distinct social groups may also act as barrier to entry or expansion.

'Herding bias', whereby individuals tend to bias their own data inputs towards those of a group they wish to be part of, can increase this sort of market segmentation effect.

E. Other Distortions to Competition

As discussed above, we may also observe data-related 'sparsity bias', whereby the availability of data is likely to be nonrandom. If an allowance is not made for such sparsity bias within RSs, then RSs are typically less likely to recommend items with less data associated with them, and this has the potential to distort competition.

The role of sparsity bias in generating popularity bias and incumbency bias has already been discussed. However sparsity bias can also lead to suppliers competing to elicit a reaction from end-users and thereby reduce data sparsity (for example, through more flamboyant marketing). It can also disproportionately benefit content that is harmful, extreme, or fraudulent. For example, it can be hard for genuine investments to compete effectively if they are having to compete with apparently far more exciting—but risky—cryptocurrency-based products.

Sparsity bias can also favour particular categories of products, thereby distorting competition between these and other categories. For example, if there is sparsity bias in movie rankings, whereby they disproportionately reflect reviews and ratings by males, then this could result in movies that appeal to men receiving higher average rankings and attracting more viewers. Movies that appeal to women may then find it disproportionately hard to gain attention and consequently investment. Likewise, sparsity bias that disfavours English language music on streaming sites may disincentivize artists from creating music that is not in English (Antal, 2020). From these examples, it is clear that sparsity bias might not only distort competition but also have a (detrimental) cultural impact.

VI. MITIGATING SYSTEMIC RECOMMENDATION BIASES

We have seen that systemic biases in RS can both worsen end-user choices and affect competition between suppliers, usually negatively, and that this can occur even where platforms wish to design their RS to be as customer-centric as possible. This raises the question of whether platforms have any ability to mitigate these systemic recommendation biases. After all, platforms do typically evaluate their RS, on an ongoing basis, with a view to ensuring that they are as effective as possible in reflecting end-user preferences. If platforms act on such evaluations to mitigate systemic biases, this should also ameliorate the associated competition concerns.

In practice, it can be hard to identify specific instances of biases. For example, if an end-user is recommended the most popular item, it can be hard to distinguish whether this is the result of 'popularity bias' or whether it is simply the most relevant recommendation for that end-user. There are, however, a variety of techniques available to platforms to evaluate RSs for a user population. Evaluations can be done using either online or offline methods. Online methods have the benefit of generating real end-user interactions and can thus be valuable for assessing how real users react to recommendations that are more diverse or novel. A/B testing is a well-known technique. However, there may be practical and ethical limitations to what can be tested in a live setting, and offline evaluations can also be helpful. They can also have benefits, for example, in terms of showing the precise impact of particular RS design choices on recommendation biases. In practice, a combination of evaluation approaches is often employed.

There has also been substantial (and ongoing) computer science research into how to enhance RS models to mitigate the biases described above. We have already discussed how changes to the choice architecture through which recommendations are presented can play an important role, in that end-user–item interactions will be more informative about true preferences if end-users have made an active choice from a list of recommendations, rather than being strongly steered towards a particular option.

Where significant systemic biases are identified, it may be possible to adapt the RS model to ameliorate these. Approaches taken to date include the following.

Hybrid models. Hybrid models incorporate a variety of the filtering systems described above, noting that they have different biases and thus their use in combination can—to some extent at least–help to ameliorate these biases.

Increasing exploration. A key method for enhancing RS models is to utilize reinforcement learning techniques (Aggarwal *et al.*, 2016). These essentially combine 'exploitation' (use of the RS model as it stands) with 'exploration' (designed to enhance the model). For example, the simple 'epsilon-greedy' method randomizes between providing the 'best' recommendation in the majority of instances and providing more experimental recommendations in the remainder of instances.

The occasional use of experimental recommendations has two important benefits. First, in some cases, the recommendation will actually be better than the 'best fit' recommendation, especially if end-users actually value diversity and 'surprise' in the recommendations they receive. Second, whether or not the recommendation pays off, the data it generates are useful for strengthening the model, and in particular for mitigating systemic biases. Moreover, explore/-exploit techniques also help reduce the tendency of data feedback loops to amplify such biases. However, such approaches are unlikely to provide a complete solution. It can be difficult to experiment too much without losing end-users or reducing end-user welfare. This can be a serious constraint in systems with a very large number of users and items.

Inclusion of additional objectives. Another approach is to include additional objectives within the RS, to be jointly optimized alongside 'relevance'. For example, some RS designs specifically include as objectives: 'diversity' (Helberger *et al.* 2018; Hamedani and Kaedi 2019; Wasilewski and Hurley, 2016); 'fairness' (Farnadi *et al.*, 2018; Mehrotra *et al.*, 2018; Abdollahpouri *et al.*, 2017); or 'serendipity' (Kotkov *et al.*, 2016, 2020; Akiyama *et al.*, 2010; Ziarani and Ravanmehr, 2021). Serendipity is a complex concept in this context as—like the explorative techniques described above—it requires the RS to recommend a product that is not necessarily relevant to what the end-user is looking for. One approach is to generate a list of relevant items whilst maximizing accuracy and then to re-rank this list. Another is to modify the accuracy-oriented algorithm, for example, by using a modified form of collaborative filtering whereby, instead of recommending items that very similar end-users chose, it deliberately recommends items chosen by end-users who are somewhat different (Said *et al.*, 2013; Tuzhilin and Adamopoulos, 2013). Finally, there are also serendipity-focused models that are not based on any common accuracy-oriented algorithms, such as Akiyama *et al.* (2010) or Kotkov *et al.* (2020).

Introducing additional objectives can be helpful in mitigating some of the biases outlined above, such as homogeneity bias or conformity bias. However, they can themselves introduce new biases and may not bring recommendations any closer to the underlying end-user preferences. For example, an attempt by Spotify to introduce a fairness criterion into its music streaming RSs was found to lead to reduced end-user satisfaction (Mehrotra *et al.*, 2018).

In this article, our core assumption is that end-user preferences are well defined and consistent. They just do not know which of the wide-range of available options best meets those preferences. However, it should be noted that this may not be true in practice, where end-users may themselves be ignorant about their own preference ranking. Moreover, in the context of data feedback loops, recommendations can themselves form and reshape end-user preferences. Whether this is positive or negative is largely subjective, but we might think it positive if an end-user discovers a great new band and negative if they become obsessively anti-vax.

This potentially dynamic nature of end-user preferences poses an additional question about RS design and, in particular, the objectives chosen for the RS: Which end-user should the RS be seeking to help, today's or tomorrow's? RSs that focus too rigidly on reflecting today's preferences are less likely to recommend new/different products and are therefore less likely to encourage new discovery and diversity. On the other hand, RSs that try to surprise the end-user, by offering novel items, may lose the end-user completely if the recommendations fall too far from the end-user's current preferences.

Complex modelling approaches. Finally, there are a variety of more complex modelling approaches, for example, using deep learning techniques or seeking to estimate propensities or utilities (Chen *et al.*, 2020). These can be used to create better estimates of end-user preferences for items where none has yet been expressed and can therefore help to mitigate a number of the biases described above. For example, random neural networks can be used to overcome popularity bias in music streaming continuation (Vall *et al.*, 2019). However, such approaches tend to be very data-intensive. Moreover, deep learning makes these systems less transparent, and therefore, it can be more difficult to identify and mitigate any (residual or additional) biases that do arise.

Overall, the key point to highlight for the purposes of this article is that—despite the valiant attempts of myriad computer scientists—these bias mitigation techniques are unlikely to remove the identified biases completely. Indeed, they may even introduce new biases, depending on precisely how the bias mitigation is done.

VII. INTRODUCING PLATFORM INTERESTS

In this article so far, we have assumed that the intention of the platform is that its RS should be customer-centric (even if this intention is not always achieved). In practice, however, platforms are commercial enterprises. As such, their primary interest is typically in generating profit for their shareholders.

In relation to RS design, it may be the case that the interests of shareholders are broadly aligned with those of end-users. After all, if RSs deliver more relevant recommendations, then end-users may be expected to buy more items, which will, in turn, generate more profit. An RS that seeks to maximize platform revenue need not harm other objectives, such as end-user satisfaction (Azaria *et al.*, 2013). Likewise, if an RS can generate more competition between upstream suppliers, then prices will tend to be lower on the platform, again attracting more end-users.

A. Tensions between Platform and End-User Interests

However, this happy coincidence of interests is not necessarily always true. There are (at least) five reasons why a platform's interests may not be entirely in line with those of end-users, and thus why a platform might not want its RS to be fully customer-centric, and may wish to include a commercial objective within the RS, such as maximizing revenue or profit.

Profit from restricted supplier competition. First, the platform may be in a position to profit from restricted supplier competition, for example, by charging a fee that extracts a share of the resulting rents. For this reason, Fletcher *et al.* (2023) show that the preferences of the platform over choice of the RS model can be the precise reverse of the preferences of end-users. In a field experiment on a video-on-demand system, Zhang *et al.* (2021) estimate the effect of using a profit-maximizing RS relative to an end-user welfare maximizing RS. Price sensitivity is found to fall dramatically for any video that is placed in a particularly prominent specific slot. Again, this implies that the profit-maximizing allocation of videos to that slot diverges from that which maximizes end-user welfare.

Favouring suppliers that confer a higher margin. Second, the platform may favour those suppliers from which it can extract a higher margin. Suppose there are two items A and B, offered by suppliers x and y, respectively, and recommending A would result in a higher level of end-user satisfaction, but recommending (and the end-user buying) B represents higher revenue for the platform. This could happen, for example, where the platform charges a higher commission to the supplier of B. In the context of streaming platforms, Bourreau and Gaudin (2021) argue that product B is likely to be recommended. Hunold *et al.* (2020) identify similar behaviour in online travel agencies. If the RS steers end-users towards suppliers that award the RS a greater margin, it could lead to suppliers competing to give RS a bigger margin. This increased cost would then feed into prices to end-users. Although not in the context of RSs, this possibility has been discussed by Armstrong and Zhou (2011), Inderst and Ottaviani (2012), and Hunold and Muthers (2017).

Vertical integration and self-preferencing. Third, the platform may be vertically integrated into supply. In this setting, a platform may have an incentive to tweak its RS to favour its own upstream product, creating a 'self-preferencing bias'. Self-preferencing is not a novel idea. What makes it different in this context is the subtlety with which self-preferencing can be achieved through RSs. It is enough if the platform only marginally tips the platform to favour some products; data feedback loops can amplify the impact of even the smallest changes to the initial conditions of the RS. With enough iterations through the RS loop, these small changes can completely alter the fate of some products and suppliers on the platform. Likewise, a self-preferencing platform may rank their own products above others for a short period and then stop. Data feedback loops mean that this initial push will be preserved and amplified in the RS, and it may be enough to tilt the competitive playing field on the platform completely. Proving any misconduct—especially where the misconduct may not even be discernible from the ongoing operation of the RS—may be very difficult in such cases.

Another form of potential self-preferencing relates to the platform's proprietary access to the data associated with the RS. Such data can be valuable for suppliers, as it can be used to predict which products are most likely to be recommended to end-users. If the platform is the only one with access to this information, it may help enable them to create products that the end-users would more likely buy than those produced by third-party suppliers. For example, in creating House of Cards, Netflix relied heavily on its own RS data in relation to design, development, and talent selection (Schrage, 2020, p. 11).

Supplier bargaining power. Fourth, some larger suppliers may have bargaining power with respect to the platform and may be able to impose contracts that influence RS recommendations. A relatively little-discussed implication of network effects is that certain suppliers can effectively become 'must have'. Without their presence on a given platform, end-users on the other side of that platform would switch to an alternative platform. This could, in turn, lead to other suppliers leaving and so on. Such critical suppliers have substantial bargaining power and can potentially utilize this to require preferential treatment by a platform's RS. This requirement can be direct, but it can also be indirect. For example, some music streaming services have minimum payment guarantees with the three major record labels, each of which has substantial bargaining power. At the margin, such minimum payment guarantees may be expected to incentivize the streaming services to favour major label music over independent music (Mariuzzo and Ormosi, 2022; Antal *et al.*, 2021).

Strategic interests. Fifth, a platform may have wider strategic reasons for distorting supplier competition. For example, a firm that offers an ecosystem with many different services within it may wish to keep end-users within its 'walled garden'. As such, even if it does not itself provide a particular product, it may be more inclined to recommend a third-party product that lies within the walled garden than one that would take end-users outside it. For example, Google's mobile search service (at one stage) gave preference in its rankings to content that was cached on Google's own AMP servers (AMP originally stood for 'accelerated mobile pages'). This may be—as Google claimed—because Google could then be sure of the download speed and quality of such content. However, it might also have reflected Google's preference to keep end-users within the Google ecosystem. Similar considerations may apply in relation to Amazon giving preference in its rankings to third-party suppliers that use its 'fulfilled by Amazon' service.[8]

It should be highlighted that not all of the above conduct would comprise illegal conduct under antitrust rules. For example, self-preferencing will typically not be anticompetitive if carried out by a nondominant platform. However, where platforms have substantial market power, it is clear that incorporating commercial objectives into RS has the potential to generate additional anticompetitive effects, over and above those discussed in previous sections.

B. RS and Advertising

The focus in this article is on RSs. However, we note that the line between RSs and advertising is not always a clear one, especially where suppliers can pay (or offer a higher margin) to gain more prominence or a higher ranking. This is the case, for example, on many hotel booking sites, whereby hotels are invited to share a larger proportion of their revenues in return for a higher ranking. In this case, the platform has to weigh two, potentially conflicting, objectives: maximizing this advertising revenue (by recommending something that may not be most

8 https://en.agcm.it/en/media/press-releases/2021/12/A528.

relevant for the end-user) and maximizing end-user utility (by recommending what is best for the end-user). The tensions raised by this particular multiobjective problem are discussed in Malthouse *et al.* (2019).

Another related issue is when the platform that deploys the RS also receives advertising revenue from the suppliers. If RS and traditional advertising are substitutes, then improved recommendations may reduce demand for personalized advertising, and a platform may therefore wish to downgrade its RS to avoid cannibalizing its own advertising revenue. A key element of the European Commission's *Google Shopping* antitrust case involved Google pushing rival price comparison sites down the organic search results, to focus end-users' attention on its own Google Shopping Box (which sells advertising).[9]

In an online environment, however, even pure advertising can take a somewhat-similar form to RS, at least where the placement of that advertising is underpinned by algorithms. Suppliers of online real estate (effectively selling end-users' eyeballs) prefer to sell that space to advertisers that will prove attractive to its end-users. This reflects the fact that much advertising online is sold on a 'cost per click' basis. If end-users are not interested in the product, and do not click, the seller of the space will not make as much money. Algorithms very similar to RSs are therefore used to predict the likelihood of clicks, and these then feed into the advertising auction process. Advertisers that are considered less likely to get clicks effectively have to pay more each click they do get.

To the extent that these sorts of algorithms exhibit similar biases to RSs, there may be similar effects arising in the online advertising market as are discussed in this article. For example, new entrants—for which there is no established click through data—may have to pay substantially more for their advertising than existing firms, further exacerbating incumbency advantages.

VIII. IMPLICATIONS FOR COMPETITION POLICY

Much of the discussion in this article relates to competition effects arising from systematic RS biases that can emerge even when those RSs are intended to be customer-centric. While it is unlikely that platforms would be sanctioned under antitrust rules in relation to the well-intentioned RS design of this sort, there may nonetheless be some important implications for antitrust from this discussion, in particular in relation to competition in supplier markets. Moreover, antitrust risk is more likely to be relevant once we allow for the possibility of divergent incentives between platforms and end-users.

In addition, even where antitrust is not necessarily applicable, there may be a role to be played in this space by pro-competition digital regulation.

A. Antitrust

Even where RS design is intended to be customer-centric, we have seen that RS biases can affect the extent of competition between products, and this can, in turn, have implications for the analysis of competition in supplier markets.

For example, if two suppliers provide very similar products, but RS design means that one is always recommended to one group of end-users, and the other is always recommended to another, then the competition between those products may in practice be limited. This can have potential implications for market definition, assessment of market power, and merger control. In addition, RSs can directly affect the ease of supplier entry and expansion. This will have clear relevance to assessing supplier market power and mergers between suppliers. A merger between

[9] General Court of the European Union, PRESS RELEASE No 197/21, Luxembourg, 10 November 2021, Judgment in Case T-612/17 Press and Information Google and Alphabet v Commission (Google Shopping), https://curia.europa.eu/jcms/upload/docs/application/pdf/2021-11/cp210197en.pdf

two suppliers is more likely to be harmful if entry and expansion are restricted due to such RS biases.However, caution is needed in such cases. RS can always be re-designed, and this can potentially change the nature of competition quite quickly. Moreover, the way in which the RSs make recommendations might be more easily gamed by a supplier that owns both products, which could be relevant to assessing merger effects.

Antitrust concerns are more likely to arise in relation to RS design where platforms and end-users have divergent interests, as discussed in the previous section. Here, the focus to date has primarily been on the issue of self-preferencing in recommendations and in particular the risk that dominant platforms may be able to leverage their market power into related markets by recommending their own services. The recent European General Court judgment on the *Google Shopping* case confirms that 'self-preferencing' can be illegal under competition law when carried out by a dominant platform.[10]

However, the discussion above highlights that identifying self-preferencing can be a complex exercise. The feedback loops underpinning RS development mean that bias in the data or model at any point in time can have a long-term impact, even if the bias is ostensibly removed at a later stage. For example, if a platform's RS favours its own products enough to make them the most popular options, then they will gain long term from the inherent popularity bias within RS, even absent further self-preferencing. Castellini *et al.* (2023) provide evidence to this in a simulated setting.

B. Pro-Competition Regulation

We noted above that any harm to supplier competition from RS biases was likely to be greatest in markets where specific platforms accounted for a large share of suppliers' potential sales. In this context, it is noteworthy that several jurisdictions are currently introducing pro-competition regulation for the largest digital platforms that act as a 'bottleneck' or 'gatekeeper' in relation to a number of critical online markets. The European Union (EU) is leading the way with the Digital Markets Act[11] (DMA), with the United Kingdom and Australia not far behind.

Such regulation is interesting in the context of this article because it can address competition harms that arise from the conduct of large digital platforms, even if such conduct was not strategically intended to cause harm.

Of particular interest within the DMA is Article 6(5): *'The gatekeeper shall not treat more favourably, in ranking and related indexing and crawling, services and products offered by the gatekeeper itself than similar services or products of a third party. The gatekeeper shall apply transparent, fair and non-discriminatory conditions to such ranking and related indexing and crawling'*. This blanket ban on self-preferencing within the RS of designated platforms almost certainly goes beyond the requirements of antitrust law. Such an obligation should help to ensure a level playing field for competition amongst the suppliers that rely on these platforms for their critical access to end-users.

Also of relevance is Article 6(1): *'The gatekeeper shall not use, in competition with business users, any data that is not publicly available that is generated or provided by those business users in the context of their use of the relevant core platform services or of the services provided together with, or in support of, the relevant core platform services, including data generated or provided by the end-users of those business users'*. This is designed to address the concern that platforms may be able use the data they collect through the operation of their RSs to leverage their position in a related market.

These DMA provisions are designed to address the situation where platform interests may diverge from those of end-users, due to the platform being vertically integrated into supply.

[10] See Footnote 8 above.
[11] Regulation (EU) 2022/1925.

However, the wider lessons of this article will be relevant in assessing compliance with these provisions. In particular, our findings suggest that it may be unrealistic to require gatekeepers to make their RSs entirely unbiased or entirely competitively neutral.

IX. WIDER POLICY IMPLICATIONS

We have seen that RSs can exhibit systemic biases that, in turn, have negative implications for competition. We discussed how these are linked to the fact that RSs are predicting end-user preferences on the basis of incomplete data. We have discussed how greater exploration, whereby RSs experiment with less obvious recommendations, can help reduce these biases over the longer term. We have seen how consumer choices may be influenced by the choice architecture, which, in turn, re-train RS models and thus influence recommendations going forward. We have also mentioned that consumer preferences may not in fact be well defined and consistent, but may rather themselves be influenced by the RS.

This set of observations raises a number of conflicting wider policy issues beyond pure competition policy. First, there is a potential tension between privacy policy and competition. Privacy policy may be expected increasingly to limit the collection and sharing of personal data.[12] But we have seen that it is the incompleteness of data that contributes to the systemic biases in RS that can, in turn, harm competition. The effectiveness of RS in providing relevant recommendations depends critically on the data inputs it is trained with. A reduction in the extent or granularity of personal data available for this task, due to privacy concerns, could have significant implications for RS going forward and consequently for competition between the suppliers that access end-users through them.

Second, and perhaps more surprisingly, the same effect could even give rise to a conflict between different competition objectives. For example, Article 5(2) in the DMA is designed to limit the combination and cross-use of personal data across different services provided by a gatekeeper platform. The intention would appear to be to promote contestability by placing rivals—with less access to such data—on a more level playing field. While this may indeed be useful for promoting inter-platform competition, the discussion in this article highlights that it could nonetheless harm intra-platform competition between suppliers.

Third, the importance of choice architecture for the long-term development of RS highlights the critical role of consumer protection policy for effective competition. However, the importance of exploration for improving RSs does raise a potential issue. Such exploration inherently involves making recommendations that are not necessarily the best in the short term, albeit the aim of improving recommendations over the longer term. It is important, therefore, that the consumer protection policy takes a long-term view when reviewing RSs and does not disincentivize such positive exploration activity on the basis that it might be misleading or deceptive in the short term under consumer law.

Fourth, the role of RS in influencing consumer preferences is not a core theme of this paper, but is clearly very important, and a focus of the new EU Digital Services Act[13] (DSA). This requires the 'very large online platforms' and 'very large search services' to analyse any systemic risk stemming from the use of their platforms and put in place effective mitigation measures. The importance of RS in this context is clear, with the DSA highlighting that '*RS play an important role in the amplification of certain messages, the viral dissemination of information and the stimulation of online behaviour*' (Recital 70) and '*providers should therefore pay particular attention on how their*

[12] The core EU Regulation relating to privacy is the 2016 *EU General Data Protection Regulation* (Regulation (EU) 2016/679). However, there are additional privacy-related requirements in other regulations. For example, Article 38 of the Digital Services Act requires that 'very large online platforms and search engines that use RS *"shall provide at least one option for each of their recommender systems which is not based on profiling"*'.

[13] Regulation (EU) 2022/2065.

services are used to disseminate or amplify misleading or deceptive content, including misinformation' (Recital 84). Finally, we discussed above how RS biases may be worsened where platforms have commercial incentives that conflict with those of end-users. Antitrust and pro-competition regulation can play a role. However, there are also a range of wider policy options, sets out below, that may be valuable in addressing the impact of such commercial incentives. Some of these are already in place, or at an advanced stage of development, but others are more novel.

A. Transparency Obligations

Transparency requirements can potentially be helpful in mitigating the impact of RS biases. This transparency could be targeted at the platform's end-users or to the suppliers that sell through it. In general, the thinking here is that, by providing more information about the RS, these users can make allowance for the way in which the RS functions, and thus the impact of any biases may be mitigated.

For example, if end-users are informed that a recommendation is partly based on the supplier having paid a fee, they might be more cautious in accepting it, albeit the impact of such a measure may depend on how explicitly it highlights a potential conflict of interest. The same may be true if consumers are warned that some recommendations may be experimental, intended to improve the quality of the RS. In the context of music streaming, greater transparency around the basis for recommendations has been found to increase end-users' confidence in those recommendations (Sinha and Swearingen, 2002). At the same time, if suppliers selling through platforms have greater understanding of the criteria underpinning RS, this may help them in ensuring that they get fairer visibility for their products. Otherwise, there is a risk that the products that sell the best are not the best products, but rather those whose suppliers better understand the ranking criteria.

The European Commission has recently introduced a suite of different transparency requirements that cover RSs. First, it now requires enhanced transparency for end-users in key consumer protection legislation. In an update to the Unfair Commercial Practices Directive[14] (UCPD), new paragraph 4a within Article 7 requires that providers of search services for third-party products (effectively, RSs) must set out the main parameters determining the ranking of those products. In addition, it is now explicit within the UCPD that undisclosed advertising and paid promotion in search results are prohibited. The update also introduces new rules around the use of consumer ratings and reviews, including in relation to the need for transparency about how they are aggregated to create aggregate rankings.[15] While these measures are targeted at protecting consumers, they may also have benefits for suppliers selling through such search services.

Second, there are relevant transparency obligations, targeted at business users, in the EU Platform to Business Regulation.[16] Article 5(5) requires that intermediation platforms '*provide a description that gives users an adequate understanding of whether—and if so, how and to what extent—the ranking mechanism takes account of: (a) the characteristics of the goods or services offered through the provider's service; (b) the relevance of those characteristics to the consumers using that service; and (c) solely as regards providers of online search engines, the design characteristics of the website used by the corporate website users*'. The European Commission has also issued guidelines on the application of these requirements.[17]

14 Directive 2005/29/EC.

15 Note that simple averages of ratings are not necessarily optimal for consumers. For example, consumers may benefit from aggregation approaches which weight ratingsor reviews based on their informational content or verifiability (Dai et al. 2018).

16 Regulation (EU) 2019/1150.

17 Official Journal of the European Union, 2019, *Commission Notice—Guidelines on ranking transparency pursuant to Regulation (EU) 2019/1150 of the European Parliament and of the Council*, 2020/C 424/01.

Third, the new EU DSA includes specific provisions relating to transparency of RS. Article 27(1) states that *'providers of online platforms that use recommender systems shall set out in their terms and conditions, in plain and intelligible language, the main parameters used in their recommender systems, as well as any options for the recipients of the service to modify or influence those main parameters'*. This must include *'(a) the criteria which are most significant in determining the information suggested to the recipient of the service; and (b) the reasons for the relative importance of those parameters'*.

Fourth, the EU is also proposing a new Artificial Intelligence (AI) Act, which includes transparency requirements.[18] Article 52(1) requires that *'Providers shall ensure that AI systems intended to interact with natural persons are designed and developed in such a way that natural persons are informed that they are interacting with an AI system, unless this is obvious from the circumstances and the context of use'*.

These various transparency requirements should prove helpful in ameliorating some of the concerns highlighted in this article. However, it is far from obvious that transparency is a panacea. Consumers and suppliers may not know what to do with the additional information they receive, even if they are aware of it.

Moreover, there are reasonable arguments for avoiding too much transparency. Platforms have concerns that too much transparency may make their RS gameable by smart third-party suppliers, such that the suppliers who end up being recommended are those that are best at gaming the RS design, not those with the most suitable products. This would not be good for either end-users or competition. Too much transparency around RS also risks making it easy for rivals to copy a platform's algorithms. This would breach their IP (trade secrets) and risk disincentivizing innovation in RS design. These latter risks do not undermine the case for transparency completely, but they may mean that the optimal level of transparency is somewhat lower than would otherwise be the case.

B. Enhanced End-User Control

Throughout this article, we have assumed that the platform is in sole charge of designing its RS and collecting the data required to train and apply it. However, there may be potential for end-users themselves to play a role here, either by inputting relevant data or by choosing some aspects of how the RS works. Many platforms effectively do elements of this already, such as Facebook and Twitter enabling its end-users some influence over the items they see in their newsfeed. However, this could go further, with regulatory support.

The new EU DSA makes a move in this direction. Article 27(3) requires that, where several options are available for RS to determine the relative order of information presented, platforms should *'make available a functionality that allows the recipient of the service to select and to modify at any time their preferred option. That functionality shall be directly and easily accessible from the specific section of the online platform's online interface where the information is being prioritised'*.

C. Professional Diligence and Fiduciary Duties

A final option is to impose on RS providers an additional duty of professional diligence in relation to their RS or fiduciary duty in relation to their end-users. We discuss above a number of ways in which biases within RS can be mitigated, albeit none of these is entirely perfect. Most good RSs will be carrying out this sort of ongoing assessment and adaptation process anyway, but others may need the motivation of a regulatory requirement. For 'very large online platforms' and 'very large search engines', the risk analysis and mitigation requirements within the EU DSA (described above) go some way in this direction.

18 COM/2021/206.

X. CONCLUSION

RSs are intended to provide relevant recommendations for end-users that reflect their preferences. However, they typically have incomplete information on these preferences, and this can lead to systemic biases in the recommendations made. These biases are important not only because they may lead to poor end-user choices but also because they may distort competition between suppliers, potentially driving higher concentration, raising barriers to entry and expansion, creating market segmentation, and reducing variety. The impact on supplier incentives may lead to higher prices and reduced innovation. The situation may be worse if a platform's own interests are not aligned with those of their end-users and more complex still if end-user preferences are not well defined and consistent but can potentially be influenced by the RS.

While we discuss some policy implications, including implications for competition policy, the primary aim of this article is to identify these key effects at a high level and, in doing so, to motivate future economic and legal research. Although RSs have been the subject of a burgeoning body of computer science literature, there has been far less analysis of the likely effects of systemic RS biases for competition in markets. Notable exceptions are Calvano *et al.* (2023), Castellini *et al.* (2023), and Fletcher *et al.* (2023). The remaining economic literature in this area focuses on the situation where market outcomes result not from systemic RS biases but rather a clear divergence of interests between the platform and its end-users.

Further economic research might include a combination of economic theory and computer simulation to examine in more detail the likely impact of biases on competitive outcomes, as well as empirical work, controlled experiments, and evaluation techniques (such as A/B testing) to enhance our understanding of the implications of RSs for supplier competition in real-world applications. Further legal work could also usefully investigate the likely impact of new laws, which may well act to limit the granularity of data available to RS, and the conflicts between various wider public policies. At a time when digital platforms are at the centre of tectonic shifts in markets and economic outcomes, such evidence will be critical for designing better informed policy and regulatory responses.

FUNDING

This work has benefited from a UKRI Trustworthy Autonomous Systems (TAS) Pump Priming grant EP/V00784X/1. The UKRI TAS Hub assembles a team from the Universities of Southampton, Nottingham and King's College London. The Hub sits at the centre of the £33M Trustworthy Autonomous Systems Programme, funded by the UKRI Strategic Priorities Fund. Amelia Fletcher's work was part-funded by EPSRC grant EP/T022493/1.

REFERENCES

Abdollahpouri, H., Burke, R. & Mobasher, B. (2017), Controlling Popularity Bias in Learning-to-Rank Recommendation, *Proceedings of the Eleventh ACM Conference on Recommender Systems*, pp. 42–46.

Abdollahpouri, H., Mansoury, M., Burke, R., Mobasher, B. 2019. The unfairness of popularity bias in recommendation, arXiv preprint arXiv:1907.13286.

Abdollah, H., Mansoury, M. 2020. Multi-Sided Exposure Bias in Recommendation, *International Workshop on Industrial Recommendation Systems with ACM KDD*, **2020**.

Aggarwal, C.C. *et al.* 2016. *Recommender Systems: The Textbook*. Springer, Vol. **1**.

Aguiar, L., Waldfogel, J. 2021. Platforms, power, and promotion: Evidence from Spotify playlists, *The Journal of Industrial Economics*, **69**(3): 653–91.

Akiyama, T., Obara, K., Tanizaki, M. 2010. Proposal and evaluation of serendipitous recommendation method using general unexpectedness, *Proceedings of the Workshop on the Practical Use of Recommender Systems, Algorithms and Technologies*, PRSAT 2010, pp. 3–10.

Anderson, A., Maystre, L., Anderson, I., Mehrotra, R. & Lalmas, M. (2020), Algorithmic Effects on the Diversity of Consumption on Spotify, *Proceedings of The Web Conference 2020*, pp. 2155–65.

Antal, D. (2020), *Feasibility Study on Promoting Slovak Music in Slovakia & Abroad*, mimeo.

Antal, D., Fletcher, A., Ormosi, P. 2021. Music streaming: Is it a level playing field? *Competition Policy International: Antitrust Chronicle*, 2(2).

Aridor, G., Gonçalves, D. 2022. Recommenders' originals: The welfare effects of the dual role of platforms as producers and recommender systems, *International Journal of Industrial Organization*, **83**: 102845.

Armstrong, M., Wright, J. 2004. Two-sided markets, competitive bottlenecks and exclusive contracts, *Economic Theory*, **32**: 353–80.

Armstrong, M., Zhou, J. 2011. Paying for prominence, *The Economic Journal*, **121**(556): F368–95.

Azaria, A., Hassidim, A., Kraus, S., Eshkol, A., Weintraub, O. & Netanely, I. (2013), Movie Recommender System for Profit Maximization, *Proceedings of the 7th ACM conference on Recommender systems*, pp. 121–128.

Belleflamme, P., Peitz, M. 2020. Ratings, reviews and recommendations. In *Handbook of Cultural Economics*, Third edn. Edward Elgar Publishing.

Bobadilla, J., Ortega, F., Hernando, A., Bernal, J. 2012. A collaborative filtering approach to mitigate the new user cold start problem, *Knowledge-Based Systems*, **26**: 225–38.

Bourreau, M., & Gaudin, G. 2022. Streaming platform and strategic recommendation bias. *Journal of Economics & Management Strategy*, **31**(1):25–47.

Brynjolfsson, E., Hu, Y., Simester, D. 2011. Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales, *Management Science*, **57**(8): 1373–86.

Calvano, E., Calzolari, G., Denicolò, V., Pastorello, S. 2023. *Artificial Intelligence, Algorithmic Recommendations and Competition*, working paper, SSRN.

Carare, O. 2012. The impact of bestseller rank on demand: Evidence from the app market, *International Economic Review*, **53**(3): 717–42.

Chaney, A. J., Stewart, B. M. & Engelhardt, B. E. (2018), How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility, *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 224–32.

Chen, J., Dong, H., Wang, X., Feng, F., Wang, M., & He, X. 2023. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems*, **41**(3): 1–39.

Ciampaglia, G.L., Nematzadeh, A., Menczer, F., Flammini, A. 2018. How algorithmic popularity bias hinders or promotes quality, *Scientific Reports*, **8**(1): 1–7.

Coyle, D., Fletcher, A., Furman, J., Marsden, P., McAuley, D. 2019. *Unlocking Digital Competition: Report of the Digital Competition Expert Panel.*

Cure, M., Hunold, M., Kesler, R., Laitenberger, U., Larrieu, T. 2022. Vertical integration of platforms and product prominence, *Quantitative Marketing and Economics*, **20**: 353–95.

Dai, W.D., Jin, G., Lee, J., Luca, M. 2018. Aggregation of consumer ratings: An application to yelp. com, *Quantitative Marketing and Economics*, **16**(3): 289–339.

De Campos, L.M., Fernández-Luna, J.M., Huete, J.F., Rueda-Morales, M.A. 2010. Combining content-based and collaborative recommendations: A hybrid approach based on Bayesian networks, *International Journal of Approximate Reasoning*, **51**(7): 785–99.

De Corniere, A., Taylor, G. 2014. Integration and search engine bias, *The Rand Journal of Economics*, **45**(3): 576–97.

De los Santos, B., Koulayev, S. 2017. Optimizing click-through in online rankings with endogenous search refinement, *Marketing Science*, **36**(4): 542–64.

Ennis, S., Lam, W. 2021. *Personalised Pricing and Disclosure: A Report for the UK Department of Business, Energy and Industrial Strategy (BEIS).*

Farnadi, G., Kouki, P., Thompson, S.K., Srinivasan, S., Getoor, L. 2018. A fairness-aware hybrid recommender system, arXiv preprint arXiv:1809.09030.

Fleder, D., Hosanagar, K. 2009. Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity, *Management Science*, **55**(5): 697–712.

Castellini, J., Fletcher, A., Ormosi, P., Savani, R. 2023. *Recommender Systems and Competition on Subscription-Based Platforms*, Working Paper, SSRN.

Fletcher, A., Ormosi, P., Savani, R., Castellini, J. 2023. *Biased Recommender Systems and Supplier Competition.* Working Paper, SSRN.

Fong, N.M. 2017. How targeting affects customer search: A field experiment, *Management Science*, **63**(7): 2353–64.

Fong, N., Zhang, Y., Luo, X., Wang, X. 2019. Targeted promotions on an e-book platform: Crowding out, heterogeneity, and opportunity costs, *Journal of Marketing Research*, **56**(2): 310–23.
Gal-Or, E., Gal-Or, M. 2005. Customized advertising via a common media distributor, *Marketing Science*, **24**(2): 241–53.
Ghose, A., Ipeirotis, P.G., Li, B. 2014. Examining the impact of ranking on consumer behavior and search engine revenue, *Management Science*, **60**(7): 1632–54.
Goldfarb, A., Tucker, C. 2019. Digital economics, *Journal of Economic Literature*, **57**(10): 3–43.
Guo, G., Zhang, J., Thalmann, D. 2014. Merging trust in collaborative filtering to alleviate data sparsity and cold start, *Knowledge-Based Systems*, **57**: 57–68.
Hamedani, E.M., Kaedi, M. 2019. Recommending the long tail items through personalized diversification, *Knowledge-Based Systems*, **164**: 348–57.
Helberger, N., Karppinen, K., D'acunto, L. 2018. Exposure diversity as a design principle for recommender systems, *Information, Communication & Society*, **21**(2): 191–207.
Hoffmann, F., Inderst, F., Ottaviani, M. 2013. Persuasion through selective disclosure: Implications for marketing, campaigning, and privacy regulation, *Management Science*, **66**(11): 4958–79.
Holtz, D., Carterette, B., Chandar, P., Nazari, Z., Cramer, H. & Aral, S. (2020), The engagementdiversity connection: Evidence from a field experiment on Spotify, *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 75–76.
Hosanagar, K., Fleder, D., Lee, D., Buja, A. 2014. Will the global village fracture into tribes? Recommender systems and their effects on consumer fragmentation, *Management Science*, **60**(4): 805–23.
Hunold, M., Kesler, R., Laitenberger, U. 2020. Rankings of online travel agents, channel pricing, and consumer protection, *Marketing Science*, **39**(1): 92–116.
Hunold, M., Muthers, J. 2017. Resale price maintenance and manufacturer competition for retail services, *The Rand Journal of Economics*, **48**(1): 3–23.
Idrissi, N., Zellou, A. 2020. A systematic literature review of sparsity issues in recommender systems, *Social Network Analysis and Mining*, **10**: 1–23.
Inderst, R., Ottaviani, M. 2012. Competition through commissions and kickbacks, *American Economic Review*, **102**(2): 780–809.
Jannach, D., Lerche, L., Kamehkhosh, I., Jugovac, M. 2015. What recommenders recommend: An analysis of recommendation biases and possible countermeasures, *User Modeling and User-Adapted Interaction*, **25**: 427–91.
Jesse, M., Jannach, D. 2021. Digital nudging with recommender systems: Survey and future directions, *Computers in Human Behavior Reports*, **3**: 100052.
Joachims, T., Granka, L., Pan, B., Hembrooke, H. & Gay, G. (2005), Accurately interpreting clickthrough data as implicit feedback, *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Association for Computing Machinery, pp. 154–61.
Kotapati, B., Mutungi, S., Newham, M., Schroeder, J., Shao, S., Wang, M. 2020. The antitrust case against Apple, Available at SSRN, 3606073.
Kotkov, D., Veijalainen, J., Wang, S. 2020. How does serendipity affect diversity in recommender systems? A serendipity-oriented greedy algorithm, *Computing*, **102**(2): 393–411.
Kotkov, D., Wang, S., Veijalainen, J. 2016. A survey of serendipity in recommender systems, *Knowledge-Based Systems*, **111**: 180–92.
Kowald, D., Schedl, M., Lex, E. 2020. The unfairness of popularity bias in music recommendation: A reproducibility study, *Advances in Information Retrieval*, **12036**: 35.
Lederrey, G. & West, R. (2018), When sheep shop: Measuring herding effects in product ratings with natural experiments, *Proceedings of the 2018 World Wide Web Conference, International World Wide Web Conferences Steering Committee*, pp. 793–802.
Lee, D., Hosanagar, K. 2021. How do product attributes and reviews moderate the impact of recommender systems through purchase stages? *Management Science*, **67**(1): 524–46.
Lee, K.H., Musolff, L. 2021. *Entry Into Two-Sided Markets Shaped by Platform-Guided Search, Technical Report, Mimeo*. Princeton University.
Liu, Y., Cao, X. & Yu, Y. (2016), Are you influenced by others when rating? Improve rating prediction by conformity modeling, *Proceedings of the 10th ACM Conference on Recommender Systems*, Association for Computing Machinery, New York, NY, pp. 269–72.
Malthouse, E.C., Hessary, Y.K., Vakeel, K.A., Burke, R., Fudurić, M. 2019. An algorithm for allocating sponsored recommendations and content: Unifying programmatic advertising and recommender systems, *Journal of Advertising*, **48**(4): 366–79.

Mandl, M., Felfernig, A., Teppan, E., Schubert, M. 2011. Consumer decision making in knowledge-based recommendation, *Journal of Intelligent Information Systems*, **37**(1): 1–22.

Mansoury, M., Abdollahpouri, H., Pechenizkiy, M., Mobasher, B. & Burke, R. (2020), Feedback loop and bias amplification in recommender systems, *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pp. 2145–48.

Mariuzzo, F., Ormosi, P. 2023. Independent v major record labels: Do they have the same streaming power (law)? *Economics Bulletin*, (Forthcoming).

Marlin, B., Zemel, R., Roweis, S., Slaney, M. 2012. *Collaborative Filtering and the Missing at Random Assumption*, arXiv preprint arXiv:1206.5267.

Mehrotra, R., McInerney, J., Bouchard, H., Lalmas, M. & 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*, pp. 2243–51.

Nagulendra, S. & Vassileva, J. (2014), Understanding and controlling the filter bubble through interactive visualization: A user study, *Proceedings of the 25th ACM Conference on Hypertext and Social Media*, pp. 107–15.

Nguyen, T. T., Hui, P.-M., Harper, F. M., Terveen, L. & Konstan, J. A. (2014), Exploring the filter bubble: The effect of using recommender systems on content diversity, *Proceedings of the 23rd International Conference on World Wide Web*, pp. 677–86.

O'Brien, M. & Keane, M. (2006), Modeling result-list searching in the World Wide Web: The role of relevance topologies and trust bias, *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, Vol. **28**, p. 1881–86.

O'Callaghan, D., Greene, D., Conway, M., Carthy, J., Cunningham, P. 2015. Down the (white) rabbit hole: The extreme right and online recommender systems, *Social Science Computer Review*, **33**(4): 459–78.

Padilla, J., Perkins, J., Piccolo, S. 2022. Self-preferencing in markets with vertically integrated gatekeeper platforms, *The Journal of Industrial Economics*, **70**(2): 371–95.

Pagano, R., Cremonesi, P., Larson, M., Hidasi, B., Tikk, D., Karatzoglou, A. & Quadrana, M. (2016), The contextual turn: From context-aware to context-driven recommender systems, *Proceedings of the 10th ACM Conference on Recommender Systems*, pp. 249–52.

Peitz, M., Sobolev, A. 2022. *Inflated Recommendations, Working Paper*, DP 336. University of Mannheim.

Said, A., Fields, B., Jain, B. J. & Albayrak, S. (2013), User-centric evaluation of a k-furthest neighbor collaborative filtering recommender algorithm, *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, pp. 1399–408.

Schoenmueller, V., Netzer, O., Stahl, F. 2020. The polarity of online reviews: Prevalence, drivers and implications, *Journal of Marketing Research*, **57**(5): 853–77.

Schrage, M. 2020. *Recommendation Engines*. MIT Press.

Sinha, R. & Swearingen, K. (2002), The role of transparency in recommender systems, *CHI'02 Extended Abstracts on Human Factors in Computing Systems*, pp. 830–31.

Sun, W., Khenissi, S., Nasraoui, O. & Shafto, P. (2019), Debiasing the human-recommender system feedback loop in collaborative filtering, *Companion Proceedings of The 2019 World Wide Web Conference*, pp. 645–51.

Tucker, C.E. 2012. The economics of advertising and privacy, *International Journal of Industrial Organization*, **30**(3): 326–9.

Tuzhilin, A., Adamopoulos, P. 2013. *Probabilistic Neighborhood Selection in Collaborative Filtering Systems*. NYU: Center for Business Analytics Working Papers (Topic).

Ursu, R.M. 2018. The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions, *Marketing Science*, **37**(4): 530–52.

Vall, A., Quadrana, M., Schedl, M., Widmer, G. 2019. Order, context and popularity bias in next-song recommendations, *International Journal of Multimedia Information Retrieval*, **8**(2): 101–13.

Waldfogel, J. 2017. How digitization has created a golden age of music, movies, books, and television, *Journal of Economic Perspectives*, **31**(3): 195–214.

Wang, T., Wang, D. 2014. Why amazon's ratings might mislead you: The story of herding effects, *Big Data*, **2**(4): 196–204.

Wasilewski, J., Hurley, N. 2016. Incorporating diversity in a learning to rank recommender system, *Proceedings of the Twenty-Ninth International Florida Artificial Intelligence Research Society Conference (FLAIRS 2016)*, 572–577.

Yang, X., Guo, Y., Liu, Y., Steck, H. 2014. A survey of collaborative filtering based social recommender systems, *Computer Communications*, **41**: 1–10.

Yao, S., Huang, B. 2017. Beyond parity: Fairness objectives for collaborative filtering, *Advances in Neural Information Processing Systems*, **30**: 2921–2930.

Zanker, M., Jessenitschnig, M. 2009. Case-studies on exploiting explicit customer requirements in recommender systems, *User Modeling and User-Adapted Interaction*, **19**(1): 133–66.

Zentner, A., Smith, M., Kaya, C. 2013. How video rental patterns change as consumers move online, *Management Science*, **59**(11): 2622–34.

Zhang, L. 2018. Intellectual property strategy and the long tail: Evidence from the recorded music industry, *Management Science*, **64**(1): 24–42.

Zhang, X., Ferreira, P., Godinho de Matos, M., Belo, R. 2021. Welfare properties of profit maximizing recommender systems: Theory and results from a randomized experiment, *MIS Quarterly*, **45**(1): 1–43.

Zhao, Z., Chen, J., Zhou, S., He, X., Cao, X., Zhang, F., Wu, W. 2021. Popularity bias is not always evil: Disentangling benign and harmful bias for recommendation, *IEEE Transactions on Knowledge and Data Engineering*, PP(99):1–13.

Zheng, Y., Wang, D.X. 2022. A survey of recommender systems with multi-objective optimization, *Neurocomputing*, **474**: 141–53.

Ziarani, R.J., Ravanmehr, R. 2021. Serendipity in recommender systems: A systematic literature review, *Journal of Computer Science and Technology*, **36**(2): 375–96.