

Discussion Paper Series – CRC TR 224

Discussion Paper No. 786
Project B 05

Foreclosure and Exploitation through Recommendation Power

Jens-Uwe Franck¹
Martin Peitz²

October 2026

¹University of Mannheim, Department of Law, and MaCCI.

²University of Mannheim, Department of Economics, and MaCCI, Email: martin.peitz@gmail.com

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

Foreclosure and Exploitation through Recommendation Power¹

Jens-Uwe Franck and Martin Peitz

17 October 2026

Abstract

Recommendation systems create value by reducing search and information costs, but they also give platforms control over visibility and effective access to users. Recommendation power can weaken competitive constraints by distorting competition among offers or by protecting the platform's gateway position against rival gateways and direct routes to users. It can also be used to extract surplus from end users or business users. We present specific theories of harm for these uses of recommendation power and illustrate them with competition cases. We then discuss how these theories of harm can be assessed empirically, accounting for how market participants respond to recommendation design.

JEL codes: L12, L13, L41, L42, L86, D83, K21.

Keywords: recommender system, recommendation bias, platform design, foreclosure, exploitation, self-preferencing, steering.

¹ This article is based on parts of the 2026 CRESSE Keynote by Martin Peitz and complements Franck and Peitz (2026). Jens-Uwe Franck: University of Mannheim, Department of Law, and MaCCI. Martin Peitz: University of Mannheim, Department of Economics, and MaCCI, martin.peitz@gmail.com. The authors gratefully acknowledge support from the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project B05). The authors have nothing else to disclose.

1. Introduction

Digital platforms reduce search and information costs by helping users select among a large number of products, services, and content. Recommendation systems perform this function by filtering, ranking, and presenting alternatives according to information about users and offers. An unfiltered set of available options is generally not a useful benchmark: selection and ordering are part of the service that the platform provides. At the same time, recommendation determines which offers users are likely to notice, compare, and choose. Where business users depend on a platform to reach those users, control over recommendation therefore affects not only user choice but also business users' effective access to demand.

Franck and Peitz (2026) refer to this conjunction of informational control and intermediation power as recommendation power. Information power arises from the ability to shape what users observe and consider; intermediation power arises where business users depend on the platform for effective access to end users. Recommendation power arises when the allocation of visibility performs both functions at once. Rankings, featured offers, defaults, labels, badges, recommendation modules, and related design choices can then influence demand while determining which business users obtain effective access to it.

The economic value of recommendation does not imply that platforms will always use this control solely to improve matching. Recommendation rules also affect platform revenues, seller incentives, and the viability of alternative routes to users. A platform may benefit from directing demand toward its own offer, toward a third party that generates a higher return, or toward sellers whose conduct supports the platform's position. It may also use prominence itself as a monetization instrument. These possibilities raise the question studied in this paper: under what conditions can the exercise of recommendation power harm competition and consumers?

Our main contribution is to organize this question around theories of harm rather than recommendation practices. By a theory of harm, we mean a hypothesis about how a particular practice may give rise to harmful effects on competition and consumers. Much of the policy debate and a growing theoretical and empirical literature have focused on self-preferencing; see Kittaka, Sato and Zenryo (2023) and Etro (2024) for surveys. The closest contribution is Motta (2023), who develops foreclosure theories of harm for self-preferencing and the degradation of interoperability by dominant platforms. We instead ask how a platform can use recommendation for foreclosure or exploitation. Self-preferencing is important, but it is neither necessary nor sufficient for identifying the relevant competitive mechanism: the same recommendation practice can serve different purposes, and similar effects can arise without favoring the platform's own offer.

The exercise of recommendation power may give rise to foreclosure or exploitation. Foreclosure can weaken competition among offers by diverting demand away from suppliers that constrain the favored offer. It can also protect the platform's position as a gateway by making competing routes to users less attractive or less viable. The latter includes practices that penalize business users for offering better terms elsewhere, discourage direct transactions, or reduce the scale of rival gateways or complementary services. Exploitative abuse can arise

when recommendation is used to steer users toward offers that generate higher returns for the platform, to increase the value of paid prominence by weakening organic alternatives, or to use information about users in ways that permit additional surplus appropriation.

We complement this analysis by considering how these theories of harm can be assessed empirically. Observed differences in ranking or prominence do not by themselves reveal the mechanism at work or establish the relevant harm. We distinguish four questions: how recommendation and the resulting visibility are measured, and how visibility affects user choice; how the relevant counterfactual should be constructed; what drives differential treatment; and whether that treatment produces the harm posited by the theory of harm. A key distinction is between counterfactuals that hold the observed market environment fixed and those that allow prices, entry, participation, and other business-user choices to adjust to a different recommendation rule. Because business users respond to the rule by which visibility is allocated, the two counterfactuals can yield different welfare assessments of the same recommendation practice, potentially even reversing the sign of the estimated welfare effect.

The remainder of the paper proceeds as follows. Section 2 develops the economic function of recommendation and introduces recommendation power. Section 3 sets out the foreclosure and exploitation theories of harm and illustrates them with competition cases. Section 4 examines how these theories of harm can be assessed empirically. Section 5 concludes.

2. Recommendation power: Economic function and definition

2.1 The economic function of recommendations

Users typically make choices without knowing all available alternatives or all relevant characteristics of those alternatives. Product quality or the value of a particular match may be uncertain before consumption, while acquiring, processing, and comparing information is costly. Search therefore consumes time and other resources, and limited attention and cognitive capacity further constrain the number of alternatives users can evaluate.

Recommendation systems reduce these frictions by selecting, filtering, and ordering the alternatives presented to users. They may rank products or sellers, highlight particular options, summarize attributes and reviews, or select a single option from a larger set. Their economic function is therefore not merely to allocate attention but to reduce information acquisition and search costs and improve matching between users and available offers. Electronic marketplaces can create value by reducing the costs of obtaining information about prices and product characteristics, while platform design can guide consumers toward products they prefer (Bakos, 1997; Dinerstein et al. 2018; Belleflamme and Peitz, 2021, ch. 6). The appropriate benchmark is consequently not an unfiltered or unordered presentation of all available alternatives. Selection and ordering are themselves part of the service.

For an economic analysis, it is useful to make the recommendation algorithm explicit. Let I_u denote the set of offers that can be shown to user u and let s_{iu} be the score assigned to offer i for that user. A simple recommendation rule selects the offer with the highest score:

$$i^R(u) \in \arg \max_{i \in I_u} s_{iu}.$$

The score can combine predicted match quality, price, ratings, delivery conditions, or any other variable used by the platform. To illustrate, let v_{iu} be the platform's estimate of user u 's gross match value for offer i and p_i its price. Consider the scoring rule

$$s_{iu} = v_{iu} - \alpha p_i,$$

with $\alpha \in [0, \infty]$. When $\alpha = 0$, the algorithm ranks by expected gross match value; when $\alpha = 1$, it ranks by consumer surplus $v_{iu} - p_i$. These cases have a welfare interpretation. Suppose the recommender ranks offers by a weighted sum of consumer and seller surplus, $\lambda(v_{iu} - p_i) + (1 - \lambda)(p_i - c_i)$, where c_i is the marginal cost of offer i and $\lambda \in [0, 1]$ (see Reimers and Waldfogel, 2023). If marginal costs are common across the offers being ranked, this is equivalent to ranking by $s_{iu} = v_{iu} - \alpha p_i$ with $\alpha = 2 - 1/\lambda$. Full weight on consumers gives $\alpha = 1$, while equal weights give $\alpha = 0$, so that ranking by gross match value maximizes total surplus. Weights favoring sellers give $\alpha < 0$, and the rule then rewards higher prices; as the weight on consumers vanishes, the ranking approaches an ordering by price alone. What appears to be a technical parameter of the scoring rule thus encodes a judgment about whose surplus the recommendation is meant to serve.

Recommendation design more generally determines which variables enter the score, how they are weighted, and how scores are translated into visibility.² This reduced-form representation is not intended to describe the engineering of a particular recommender system. It makes explicit the economic choice embedded in the algorithm: the criterion according to which scarce visibility is allocated. The weight α may be a decision variable of the platform, and the platform's preferred value will depend on how it benefits from the transactions it intermediates, a point taken up in Section 3.2.1.

Recommendation may also generate information for users rather than merely reduce their costs of acquiring it. Che and Hörner (2018) study a recommender through which current users' experiences generate information for later users. To induce experimentation, even a recommender maximizing user welfare may sometimes recommend a product more often than current information alone would warrant. Thus a departure from the myopically best predicted match need not reflect a conflict between the platform and its users; it may arise from the value of learning for future users.

² A non-neutral scoring rule can be written as $s_{iu} = v_{iu} - \alpha p_i + \theta' z_i$, where z_i collects characteristics of offer i that enter the score but do not affect v_{iu} , such as ownership by the platform, the commission paid by the seller, the seller's price on competing channels, or its use of a platform service. An empirical test of recommendation bias then asks whether $\theta \neq 0$ conditional on price and on observable determinants of match value. Reimers and Waldfogel (2023), for example, add a term for platform-favored products to the ranking index; see Sections 4.2 and 4.3.

The choice of recommendation criterion affects market equilibrium and not only the quality of the match. Calvano et al. (2023) and Rhodes and Zhou (2026) study personalized recommendations in a search market with competing sellers. Calvano et al. (2023) run simulations of a recommender trained on sparse and noisy feedback. Recommendations create personalized prominence: each consumer begins searching from the product with the highest predicted match value, with price playing no role in the ranking. This improves matching and economizes on search, but it also raises equilibrium prices: consumers who first see a product that matches them well are less likely to search further, which makes the demand faced by each seller less price-elastic.

In Rhodes and Zhou (2026) only those consumers who disclose their data can receive a personalized recommendation and thereby economize on costly search. When the recommender selects the product with the highest match value ($\alpha = 0$), consumers who receive a recommendation become less likely to inspect non-recommended products. Sellers anticipate this reduction in price sensitivity and may raise prices. Rhodes and Zhou (2026) also consider recommendations based instead on consumer surplus ($\alpha = 1$), so that price enters the recommendation criterion. Competition for recommendation then makes demand more price-sensitive and can lead to lower equilibrium prices.

The mechanism is not confined to personalized recommendation: Dinerstein et al. (2018) study eBay's move from search results that ranked loosely filtered listings by relevance to a two-stage design in which the product is identified first and competing sellers of it are then ranked by price. Transaction prices fell, and their estimated model attributes the fall to the change in ranking criterion rather than to consumers seeing more offers.

Recommendation design can therefore affect welfare not only through the quality of the match for given offers, but also through sellers' responses to the rule by which visibility is allocated. These studies suggest two natural ways to evaluate a recommendation rule. A supply-taken-as-given assessment asks how well the rule matches users to a given set of offers at given prices and qualities. An equilibrium assessment also accounts for how sellers adjust prices, quality, entry, and participation in response to the rule. The two need not agree. Taking prices as given, ranking by gross match value ($\alpha = 0$) maximizes total surplus, but it can raise equilibrium prices; placing weight on price ($\alpha > 0$) sacrifices some match value at given prices but can intensify price competition. The equilibrium perspective is central both to the theories of harm in Section 3 and to the counterfactuals in Section 4.

2.2 Information, intermediation, and recommendation power

Precisely because users rely on recommendation to reduce search and information costs, control over recommendation gives the intermediary information power: the ability to shape the information environment in which users choose. The intermediary determines which alternatives are discovered, which are made salient, how they are ordered and compared, and which ratings, labels, or other signals accompany them. The extent of this power depends on user reliance. If users independently inspect all relevant alternatives, changes in ranking or prominence have little effect. When users instead rely on the intermediary for discovery and

evaluation, an offer that is not displayed may not be discovered at all, while differences in prominence can translate into differences in demand.³

For a platform connecting different market sides, the same allocation of visibility also determines business users' access to consumers. Intermediation power describes the dependence of business users on the intermediary as a route to end users. What matters is effective rather than merely formal access. A seller may be listed on a marketplace yet receive little effective access if it rarely appears in relevant searches, is excluded from a prominent recommendation module or Buy Box, or lacks signals that strongly affect user choice. Intermediation power is greater when demand is concentrated on the platform, alternative routes to users are weak, and business users therefore depend on the platform to reach customers at scale. It does not require literal single-homing: a platform can constitute an important route to demand even when users and business users also use alternatives.⁴

Recommendation power arises when these two elements coincide. We use recommendation power to refer to the platform's ability to steer user choice through the allocation of visibility when that visibility is also an important determinant of business users' effective access to users (Franck and Peitz, 2026). Visibility is to be understood broadly. Recommendation affects choice not only through which offers users consider but also through how they evaluate them: an offer may lose demand because it carries an unfavorable label or lacks a badge that competing offers carry. Recommendation power is therefore not confined to ranked lists. It can operate through recommendation modules, Buy Boxes, badges, ratings, defaults, featured placements, feeds, notifications, or single-result interfaces whenever these design choices materially affect which alternatives users discover and how they evaluate them.

Recommendation power should not be understood as a distinct category of market power. An infringement based on abuse typically presupposes dominance or another legally defined position of economic power that affords the undertaking substantial freedom from competitive constraints and the potential to weaken them further through unilateral conduct. Control over recommendation is one of the instruments through which such a position is exercised. It can also strengthen the position: by making alternative routes to users less viable, recommendation design can increase business users' dependence on the platform (Section 3.1.2). Relevant evidence therefore concerns both sides: how far users rely on the platform's recommendations rather than searching independently (see Section 4.1), and how far business users depend on the platform for access to demand, given the alternative routes available to them.

Recommendation power is not inherently harmful. The ability to select among alternatives and direct users toward promising matches is what allows recommendation systems to reduce search frictions in the first place. The same control over visibility, however, gives the platform an instrument with which to pursue objectives other than improving users' matches. A change

³ The terminology follows Franck and Peitz (2026), who develop the concept of recommendation power from information power and intermediation power; on information power in search, see also Peitz and Schweitzer (2016).

⁴ For the underlying account of intermediation power and effective access, see Franck and Peitz (2026), drawing on Schweitzer et al. (2018), Schweitzer and Welker (2019), and Crémer, de Montjoye, and Schweitzer (2019).

in visibility reallocates demand among business users while also changing what end users choose. A demoted seller may thus lose a large share of its sales while a favored seller gains them, and users may be better or worse off depending on how the offer they now choose compares with the one they would otherwise have chosen. Changes in the sales of individual business users are therefore not by themselves evidence of harm to competition or to users. The next section considers when control over visibility can be used to weaken competitive constraints or to extract surplus.

3. Recommendation power: Theories of harm

The potential harm from a recommendation practice depends on the economic mechanism at issue, which may involve foreclosure or exploitation. It may be used to weaken the competitive constraints imposed by actual or potential rivals and thereby protect, strengthen, or extend market power, or it may be used to exploit existing market power by extracting surplus from end users or business users. We begin with foreclosure because it connects most directly to established theories of exclusionary conduct, before turning to exploitative uses that do not require the exclusion or weakening of rivals. The same practice, such as self-preferencing, can support either theory; what distinguishes them is the causal route to harm.

3.1 Foreclosure through recommendation power

3.1.1 Preferential recommendation and competition among offers

A first foreclosure mechanism arises when a platform uses recommendation to favor one offer over competing offers and the resulting reallocation of demand weakens the competitive constraint imposed by the disfavored suppliers. The canonical case is self-preferencing by a vertically integrated platform that both intermediates transactions and supplies a competing product or service.⁵ By controlling ranking, prominence, recommendation labels, or similar forms of effective visibility, the platform can shift demand toward its own offer and away from rival offers. The relevant foreclosure concern is not preferential treatment as such, but the resulting weakening of competition faced by the favored offer.

An integrated platform may have an incentive to favor its own offer because it earns the full downstream margin on its own sales but only an intermediation fee on third-party sales. If the platform could capture rival suppliers' profits through the terms on which it grants access, it would have little reason to divert demand away from them. In digital markets, however, the platform's ability to extract such profits is often limited: organic visibility is typically not priced, and commissions are often uniform across sellers, so the platform may earn more by shifting demand to its own offer. Whether such preferential treatment gives rise to foreclosure, however, depends on its effect on the competitive constraint imposed by rival suppliers. Such foreclosure may be partial rather than complete.

⁵ For empirical evidence on preferential treatment in recommendations and rankings, see Chen and Tsai (2024), Farronato, Fradkin and MacKay (2023), and Jürgensmeier and Skiera (2025).

The Google Shopping case illustrates this foreclosure mechanism.⁶ Google gave its own comparison-shopping service prominent placement and enhanced presentation in general search results, while rival services appeared only as generic search results and were subject to demotion algorithms. The European Commission's theory was that this differential treatment diverted traffic from rival comparison-shopping services to Google's own service, thereby weakening their ability to compete effectively.

The case illustrates why recommendation power matters for foreclosure even when rivals retain formal access to users: differences in effective visibility can reallocate demand between competing offers. In a model inspired by the case, Motta (2023) shows that demoting rivals can reduce consumer surplus even when rivals can reach part of the market without the platform, so that access to the platform need not be indispensable.⁷

The Coupang case illustrates that preferential recommendation need not operate only through ranking position.⁸ The Korean competition authority found that Coupang favored private-label and directly purchased products not only by reserving high-ranking positions for them but also by having employees post reviews and award high ratings to its private label products. According to the authority, these practices increased the visibility and perceived quality of Coupang's own products, misleading consumers and diverting demand away from competing third-party sellers. The case thus illustrates how preferential recommendation can operate through both ranking and trust signals such as ratings and reviews.

Preferential treatment of an own offer, however, is not necessarily anticompetitive. Its effects depend on how sellers adjust prices, quality, or participation in response to the platform's ranking policy. Hagiu, Teh and Wright (2022), for example, study platforms that both host third-party sellers and offer their own products alongside them. They show that the welfare effects of self-preferencing depend on equilibrium responses by the platform and third-party sellers.⁹ An assessment that takes supply as given and compares the visibility of existing offers while holding seller behavior fixed can therefore be misleading.¹⁰ Self-preferencing may, for instance, intensify competition by inducing third-party sellers to respond with lower prices or higher quality. The effects of preferential treatment must also be distinguished from

⁶ European Commission, 27.6.2017, Case AT.39740, Google Search (Shopping), C(2017) 4444 final; General Court, 10.11.2021, Case T-612/17, Google and Alphabet v European Commission, ECLI:EU:T:2021:763; ECJ, 10.9.2024, Case C-48/22 P, Google and Alphabet v Commission (Google Shopping), ECLI:EU:C:2024:726.

⁷ Motta (2023) also points to a further possible rationale: by keeping users from becoming loyal to rival comparison-shopping services, which they might then access directly, Google could reduce the risk of being bypassed. The same conduct would then also support the gateway-protection theory of Section 3.1.2.

⁸ Korea Fair Trade Commission (KFTC), 5.8.2024, Decision No. 2024-284, Case Nos 2021Seogam1429 and 2022Seogam0713, Coupang and CPLB; see also KFTC, 13.6.2024, Press Release, "KFTC Imposes Severe Sanctions on Coupang for Consumer Deception". The KFTC decided the case under the prohibition of unfairly inducing a competitor's customers (Article 45 MRFTA) rather than under the prohibition of abuse of a dominant position (Article 5 MRFTA).

⁹ Relatedly, Etro (2023) endogenizes seller entry and platform commissions, showing how platform entry affects participation and welfare through these margins. Zennyo (2022) shows that self-preferencing can induce the platform to lower its commission in response to reduced third-party sales, which may lower seller prices and support consumer and seller participation.

¹⁰ For empirical analyses of Amazon that account for responses of seller prices and entry (Lee and Musolf, 2026) or of seller prices and consumer search (Lam, 2022) to a change in the recommendation rule, see Section 4.4.

those of the platform's entry with its own offer, which may increase variety or intensify price competition.

Self-preferencing is nevertheless only a special case. The foreclosure mechanism does not depend on ownership of the favored offer: preferential recommendation may also benefit an affiliated firm or a particular independent third-party supplier. What matters is the competitive relationship between the favored and disfavored offers and the effect of preferential recommendation on that relationship. By contrast, the fact that a platform favors a high-commission third-party offer does not by itself establish foreclosure.¹¹ It becomes a foreclosure theory only where the differential treatment weakens rivals in a way that reduces the competitive constraints faced by the favored supplier.

The effect of preferential recommendation need not be limited to the immediate reallocation of demand. Where visibility responds to seller choices, the recommendation rule also affects sellers' incentives to compete. A platform that designs its recommendation rule to reward lower prices or higher quality may thereby induce sellers to make more attractive offers. Conversely, if a favored offer receives preferential visibility regardless of competing sellers' pricing or quality decisions, disfavored sellers may obtain a lower return from cutting prices or improving quality, which may soften price or quality competition between the favored and disfavored offers. The competitive effect of preferential recommendation therefore depends in part on sellers' equilibrium responses to the recommendation rule.

3.1.2 Recommendation power and protection of the platform's gateway

A second foreclosure mechanism arises when recommendation is used to protect the platform against alternative routes through which business users and end users could interact. Unlike the mechanism described in Section 3.1.1, the primary concern is not competition between individual offers but competition between alternative routes through which business users can reach end users. Recommendation may thus be used to disadvantage direct sales channels, rival intermediaries, or other routes that threaten the platform's role in bringing the two user groups together. Such foreclosure does not require the favored offer to be supplied by the platform itself. Recommendation can also favor particular third parties because their conduct or future development poses less of a threat to the platform's gateway position.

The most immediate mechanism arises when visibility is conditioned on a business user's conduct outside the platform. A seller that relies on the platform to reach customers may also sell through its own website or a competing intermediary, possibly at a lower price. The platform can discourage such bypass without prohibiting it contractually by reducing the seller's visibility when it offers more attractive terms elsewhere. The seller then faces an opportunity cost of using the alternative channel: a lower off-platform price may increase direct sales but reduce the traffic it receives from the platform. Recommendation can in this way operate as a non-contractual substitute for a price parity clause or no-steering obligation. Formal multi-homing remains possible, but becomes less effective if commercially valuable visibility on the platform depends on the seller's conduct on competing channels. Hagiu and

¹¹ If the consequence is simply that users are steered toward an offer from which the platform earns more, the natural theory is the exploitative one developed in Section 3.2.1.

Wright (2024) develop this trade-off explicitly: among the instruments through which a marketplace can limit leakage, they consider hiding sellers that induce consumers to transact outside the marketplace.

Importantly, the recommendation system can discipline off-platform pricing even when off-platform prices do not enter the ranking rule. Bar-Isaac and Shelegia (2026) consider a platform that conditions future prominence on past on-platform sales, which sellers can raise by increasing the off-platform price. Because current sales affect future prominence, the platform can use the recommendation system as an indirect instrument to reduce or even eliminate off-platform sales. The same logic applies to related sales-based metrics such as the conversion rate (on-platform sales per view): a seller whose product is frequently discovered on the platform but purchased more cheaply through a direct channel exhibits a lower conversion rate and is therefore ranked lower.

Empirical evidence from hotel-booking platforms is consistent with rankings being used to discipline off-platform pricing, although it does not identify the precise mechanism through which this occurs. Hunold, Kesler and Laitenberger (2020) find that hotels obtain better positions in an online travel agent's search results when the prices they charge on competing channels are higher than the prices charged on the platform. They interpret this pattern as consistent with online travel agents using rankings to discourage aggressive off-platform pricing.

The Bundeskartellamt's decision on Amazon's price-control mechanisms concerns visibility penalties tied to sellers' prices; one of the authority's theories of harm relates to competition outside Amazon's marketplace.¹² Amazon removed offers from its marketplace or excluded them from the Buy Box when their prices exceeded thresholds set by Amazon, which could take account of prices observed outside Amazon's marketplace. According to the authority, the threat of losing visibility induced sellers to adjust their prices on Amazon. The authority further argued that the resulting tendency to match lower off-platform prices on Amazon could reduce consumers' incentives to shop elsewhere and discourage rival online retailers from competing through lower prices.

Gateway protection can go further. A platform may care not only about whether a business user bypasses it today, but also about which business users obtain the traffic, revenues, and customer relationships generated by recommendation. Some sellers or complementors may remain dependent on the platform, whereas others may develop a direct customer base, support a competing intermediary, or eventually become a competitive threat to the platform itself. Recommendation can then be used to shift demand away from the latter and toward business users whose expansion is less likely to undermine the platform's gateway position. The relevant characteristic determining preferential treatment is therefore not necessarily ownership or current platform revenue; it can be the business user's strategic relationship to the platform.

¹² Bundeskartellamt, 4.2.2026, Case No. B2-73/20, *Amazon Marketplace—Price-Control Mechanisms*, paras. 36–81 (price-control mechanisms), 365–379 (pressure on sellers' pricing), and 474–493 (competition with online retailers outside Amazon). Amazon has challenged the decision before the German Federal Court of Justice.

This possibility gives recommendation power a natural dynamic-foreclosure interpretation. In Carlton and Waldman (2002), success in a complementary market can facilitate subsequent entry into the monopolized primary market, giving the incumbent an incentive to restrict that success today through tying. Fumagalli and Motta (2020) show a related mechanism operating through refusal to supply: by depriving an efficient downstream rival of the profits it needs to succeed, a vertically integrated incumbent can deter later entry upstream. Motta (2023) adapts this logic to digital platforms and emphasizes that vertical integration into the complementary market is not necessary: if the success of an app or merchant today provides the user base, data, or profits needed to launch a substitute platform tomorrow, the incumbent may have an incentive to limit its current expansion. Recommendation provides a particularly flexible instrument for doing so. Rather than denying access, the platform can leave the potentially threatening business user on the platform but allocate less visibility to it, while directing demand toward third parties whose success is unlikely to support future entry into the platform’s core market.

Differences in business users’ ability to develop an independent route to consumers may arise for several reasons. One is reputation. Choi and Peitz (2018) show that a supplier whose quality is not yet established can obtain market access by relying on the reputation of an established downstream firm; the intermediary effectively certifies the supplier’s quality.¹³ In a platform setting, such heterogeneity implies that traffic need not have the same strategic value for all sellers. A seller with an established independent reputation or a credible path to direct distribution may be better placed to convert current platform success into a viable outside channel than a seller whose market access continues to depend on the platform. More generally, the platform may have stronger incentives to restrict the expansion of business users that can support or become an alternative gateway than of otherwise comparable users that lack this capacity.

The Korean Google Play case provides a particularly useful illustration because the favored products were supplied by independent third parties. The KFTC, the Korean competition authority, found that Google offered game developers prominent placement and other promotional support on Google Play on condition that they release important games exclusively on Google Play rather than also on One Store.¹⁴ The foreclosure theory does not turn on Google favoring its own game. Preferential visibility made exclusivity more attractive to game developers, thereby depriving a competing app store of important content.¹⁵ The

¹³ While Choi and Peitz (2018) do not study recommendation, their analysis is relevant here because it provides a reason why business users may differ in their dependence on an established intermediary for access to consumers: some can establish an independent reputation, while others benefit from “renting” the reputation of the intermediary.

¹⁴ Korea Fair Trade Commission (KFTC), 20.7.2023, Decision No. 2023-103, Case No. 2020Jigam1764, *Google App Store*, paras. 55–76, 96–214; see also KFTC, 11.4.2023, Press Release, “KFTC Sanctions Google for Anticompetitive Conduct in App Market Blocking Developers from Releasing Games on Competing Platform”.

¹⁵ Carroni, Madio and Shekhar (2024) show how exclusivity over important content affects competition between platforms. In their model, no party is vertically integrated: two platforms compete for a leading creator and a fringe of complementors, exclusivity over the leading creator draws consumers to the favored platform, and cross-group externalities make the shift self-reinforcing as complementors follow. There, exclusivity is secured through the platform’s contractual terms; in the Google Play case, preferential visibility served as an inducement to accept exclusivity. Exclusivity is not unambiguously harmful in their setting because the entry cascade

recommendation advantage was thus strategically valuable not primarily because it altered competition between games within Google Play, but because it helped Google secure exclusive content and thereby weaken competition between app stores.

A related mechanism operates when recommendation is tied to the use of complementary platform services. Conditioning visibility on the adoption of such a service makes recommendation function as a virtual tie: the business user is not required to take the service, but declining it carries a visibility penalty.¹⁶ In the Italian Amazon FBA case, the competition authority found that sellers using Fulfilment by Amazon (FBA) obtained advantages including Prime eligibility and a greater probability of selection for the Buy Box. According to the authority, conditioning these visibility advantages on FBA diverted demand away from rival logistics providers and impeded their development. The resulting dependence on FBA also discouraged sellers from offering their products on competing marketplaces, thereby weakening competition between marketplaces.¹⁷ Motta (2023) interprets the case as a customer-foreclosure mechanism: diverting sellers toward FBA denies business to competing logistics providers, which can raise their costs and make rival marketplaces less attractive. Again, the recommendation lever matters because access to visibility on the incumbent platform affects the scale available to an alternative route.

These dynamic effects are especially important where the quality of an alternative gateway depends on scale. Traffic today may affect tomorrow's data and service quality; participation on one side may attract participation on the other; and logistics density may affect delivery cost and speed. A relatively small visibility disadvantage can therefore have consequences beyond the transactions initially diverted. Athey and Scott Morton (2022) describe a feedback mechanism in their analysis of platform annexation, where control over tools facilitating multi-homing can steer users toward the incumbent and away from rivals, weakening the rivals as their scale declines. Hagiu and Wright (2023) model the role of data directly. Whether learning from customer data confers a durable advantage depends on how quickly the returns to additional data diminish, on how much data has already accumulated, and on whether one firm learns more effectively than the other. A rival deprived of traffic accumulates less data and improves more slowly, but this need not translate into a lasting disadvantage. Whether such scale-based feedback gives rise to foreclosure in a particular case is therefore an empirical question: it depends on whether reduced traffic or participation today materially impairs the alternative route's future ability to compete.

benefits complementors on the favored platform, and consumers can gain where cross-group externalities are strong. Whether preferential visibility used to secure exclusivity raises a competition concern therefore depends on conditions and cannot be inferred from the exclusivity arrangement itself.

¹⁶ Choi, Jeon and Whinston (2026) develop a leverage theory for contractual ties of this kind. Where the platform cannot fully extract the surplus that business users derive from access, those with most to gain accept the condition whatever others do, and in an adjacent market with network effects this guaranteed base degrades rivals by depriving them of scale. Rivals are weakened rather than driven from the market, and the mechanism requires no precommitment by the platform.

¹⁷ Autorità Garante della Concorrenza e del Mercato (AGCM), 30.11.2021, Case A528 – *FBA Amazon*, Decision No. 29925.

3.2 Exploitative uses of recommendation power

Section 3.1 concerned recommendation practices that weaken competitive constraints. Recommendation power may instead be used to appropriate additional surplus from end users or business users under existing competitive constraints. We refer to such uses as exploitation; whether such conduct satisfies the legal threshold for abuse is a separate question (see, e.g., Balasingham and D’Amico, 2025).

3.2.1 Profit-driven steering and surplus extraction

A first exploitative use of recommendation power arises when the platform steers users toward offers that yield it a higher payoff. The favored offer may be supplied by a third party or by the platform itself, and the payoff advantage may arise from commissions, lead payments, royalties, transaction value, or the platform’s own downstream margin.

Such steering need not reduce user welfare. The effect depends on the alignment between the platform’s payoff and user surplus. De Cornière and Taylor (2019) make this point in a model in which consumers rely on an intermediary when choosing among sellers. If the payoff of the intermediary is aligned with consumer surplus, favoring the offer that is more profitable for the intermediary can induce the favored seller to make a better offer and may benefit consumers. If platform payoff and consumer surplus are instead in conflict, the favored seller provides lower utility and biased intermediation harms consumers. Thus, a recommendation rule that places positive weight on platform revenue is not problematic as such. The relevant question is whether the platform’s commercial criterion favors offers that also generate high user surplus, or instead favors more profitable offers at users’ expense.

Self-preferencing by a vertically integrated platform is one form this misalignment takes. When a vertically integrated platform earns the full margin on its own offer but only an intermediation fee on a third-party transaction, it may prefer its own offer even if, at given prices and qualities, the third-party offer gives the user greater surplus. In this case self-preferencing supports an exploitative theory of harm if the relevant effect is the direct diversion of users toward the less attractive but more profitable offer.¹⁸ The exploitative theory does not require a weakening of the competitive constraint imposed by disfavored suppliers; the foreclosure theory does. Relatedly, in ad-funded business models, a search engine that sells advertising alongside the content providers it links to may direct users in ways that protect its own advertising revenue rather than serve them best.¹⁹

¹⁸ Bougette, Budzinski and Marty (2022) similarly emphasize that self-preferencing may give rise to exploitative concerns.

¹⁹ de Cornière and Taylor (2014) and Burguet, Caminal and Ellman (2015) develop a version of this logic for search engines, where the platform’s return differs across destinations because it competes with content providers for advertising revenue. In de Cornière and Taylor (2014), ads on the search engine and on content providers are imperfect substitutes, so directing users to a content provider depresses the value of the search engine’s own advertising; users dislike advertising, and integration reduces advertising levels even as it increases bias, so users may gain on balance. In Burguet, Caminal and Ellman (2015), where the search engine’s sponsored search advertising competes with the display advertising sold by content providers, producer margins are misaligned with user value, and the search engine may distort either the reliability of product search or that of content search, generically distorting only one; integration with all content providers internalizes the externality and

The same mechanism operates among third-party offers. With heterogeneous commissions, lead payments or royalties, the platform earns different amounts from otherwise comparable transactions. A marketplace may therefore favor sellers paying higher commissions; a lead-generation intermediary may benefit from directing users through options that generate additional paid visits; and a streaming platform may prefer content carrying lower royalty payments. In each case, the platform's recommendation criterion departs from user value because the private return attached to an option differs across offers. Hagiu and Jullien (2011) show more generally why an intermediary may deliberately divert search when doing so raises revenue per consumer visit or affects sellers' strategic choices. Bourreau and Gaudin (2022) develop the corresponding cost-side logic for streaming platforms, where differences in royalty payments can induce strategic recommendation bias.

Importantly, differential payments are not necessary. Suppose that the platform charges all sellers the same ad valorem commission. Its revenue from a transaction then increases with the transaction price as well as with the probability of purchase. Holding conversion fixed, the platform has an incentive to direct users toward higher-priced products or larger baskets. Again, this need not reduce consumer surplus: a higher price may reflect higher quality or a better match. But when expenditure and consumer surplus are imperfectly aligned, a revenue-based recommendation rule can favor an offer that generates a larger commission even though another offer gives the user higher net value. In terms of the scoring rule of Section 2.1, such a rule places less weight on price than ranking by consumer surplus does ($\alpha < 1$) and may even reward higher prices for a given match value ($\alpha < 0$). Uniform commissions therefore eliminate one source of differential treatment while leaving intact the incentive to steer through differences in transaction value.

The same basic misalignment can arise when the platform uses conversion as a proxy for recommendation quality. Conversion is often informative: products that convert well may simply be good matches. Yet conversion can also reflect dimensions that are privately valuable to the platform but not to the user. An offer that users frequently discover on the platform but purchase more cheaply through a direct channel exhibits a lower conversion rate. Ranking by conversion then penalizes precisely those offers for which users have attractive outside purchasing opportunities. Where the main effect is to reduce users' access to attractive offers, the mechanism has an exploitative component; where it protects the platform against direct channels or rival intermediaries, it also gives rise to the gateway-foreclosure mechanism discussed in Section 3.1.2.

Conversion-based steering can be even more problematic when consumers make systematic mistakes. Offers with hidden charges, drip pricing, or costs that consumers underestimate may convert particularly well precisely because users do not fully anticipate what they will pay or receive. A platform that ranks mechanically by conversion or revenue may then amplify rather than correct these mistakes. Heidhues, Köster and Kőszegi (2023) analyze steering by intermediaries that possess information about consumers' tendencies and show how steering

improves reliability, whereas integration with only a subset gives the search engine a reason to divert users away from non-affiliated sites.

can exploit predictable consumer errors. In this setting the recommendation system need not mismeasure its commercial objective: the problem is that a high conversion probability itself partly reflects an exploitable consumer mistake.

Recommendation affects not only which existing offer receives demand but also what sellers choose in order to obtain visibility. If higher commissions increase the probability of being recommended, sellers compete partly through payments to the platform rather than through lower prices or higher quality. Armstrong and Zhou (2011) show how paying for prominence can shape competition among sellers, while Hagiu and Jullien (2011) emphasize that search diversion can influence sellers' strategic variables, including prices. Higher platform payments may subsequently be passed through to users. Conversely, recommendation design that rewards low prices or high quality can discipline sellers. The effects on end users and business users therefore depend in part on how prices and other seller choices respond to the recommendation rule.

Sellers also compete for visibility through prices rather than through payments to the platform. Zhou and Zou (2023) show that sellers adjust prices to compete for personalized recommendations and that, when the platform predicts consumer preferences sufficiently accurately, it can benefit from excluding price information from its recommendation decision (which corresponds to $\alpha = 0$ in the scoring rule of Section 2.1) because doing so prevents competition for recommendation. More generally, seller competition is an important consideration in platform strategy (Belleflamme and Peitz, 2019). In a dynamic setting, Johnson, Rhodes and Wildenbeest (2023) show that demand-steering rules that reward price cuts with additional exposure can disrupt collusive outcomes. Although their result concerns a competition-enhancing policy, it illustrates the more general mechanism: by determining how seller conduct translates into visibility, recommendation design can strengthen or soften competition among business users. The mechanism does not require coordination among sellers, although in repeated interaction recommendation design can also affect whether collusion among sellers can be sustained.

Marstaller (2026) uses the scoring rule of Section 2.1 and studies a monopoly platform that hosts sellers with a common, strictly positive marginal cost and chooses the price sensitivity of its recommendations, α , jointly with its commission rate. Unconstrained, it puts all weight on price, recommending the cheapest product ($\alpha = \infty$). Competition among sellers then drives their markups to zero, and the platform sets its ad valorem commission, which sellers pass through to consumers, so that the price equals the monopoly price. Intense on-platform competition and substantial platform extraction can thus go together. When the commission is capped, the platform makes recommendations less price-sensitive, so that consumer surplus rises by less than an analysis holding the recommendation rule fixed would predict. This can also be read as a statement about uniform commissions: where a single rate applies across a broad set of products, it binds on those for which the platform would otherwise charge more, and on those products the platform can be expected to place less weight on price, leaving sellers with higher markups.

An exploitative theory of harm can therefore operate through both the allocation of demand and the incentives created by the recommendation rule. The platform may increase the surplus it captures by directing users toward offers that generate a higher return while providing lower user surplus, or by choosing a recommendation rule that softens seller competition and allows it to capture part of the resulting seller rents. The source of the platform's additional return—its own margin, commissions, royalty savings or higher transaction value or conversion rates—is secondary. For exploitation, the central question is whether recommendation increases platform surplus at the expense of end users or business users, either directly through the allocation of demand or through its effect on seller conduct. Because the same recommendation rule may affect the two groups differently, as well as users within each group, the answer may differ across users (see Section 4.4).

3.2.2 Monetizing prominence and degrading the organic alternative

Sponsored prominence is a particular form of profit-driven recommendation in which business users pay for preferential visibility. Selling prominence is not necessarily harmful. When willingness to pay is positively related to relevance, paid placement can help sort offers by relevance (Chen and He, 2011). Conversely, a platform may set the price of sponsored listings low enough to admit less relevant offers, which raises consumers' search costs and softens competition among sellers (Eliaz and Spiegler, 2011). Whether paid placement improves matching depends in part on how the sponsored positions themselves are allocated. Athey and Ellison (2011) show this in a position auction with rational sequential search: reserve prices screen out advertisers unlikely to meet a consumer's need, so auction design determines how well the paid positions serve end users.

Auction design is itself only one of the instruments through which a platform can monetize visibility. Bar-Isaac and Shelegia (2025) study a platform that earns an ad valorem fee and chooses whether to steer by algorithm or by auctioning the privileged position. In a single product market with an endogenous fee the two are equivalent; they diverge once a common fee must be applied across heterogeneous product markets, with auctions performing better where demand or cost conditions vary sufficiently, and algorithms performing better where sellers have substantial market power or consumers differ in how far their choices are constrained by the privileged position. The choice matters for what sellers compete on: where prominence is auctioned, sellers bid for the position rather than competing for it through their product-market prices.

Vertical integration can interact with the monetization of sponsored prominence. Long and Amaldoss (2024) study a platform that sells a private-label product while also allocating a prominent sponsored position. The platform may have an incentive to favor its private label in that position, and such self-preferencing reduces consumer welfare in their model even though the private-label product is priced below competing third-party products. Some third-party sellers may nevertheless benefit. The analysis also shows that private-label sales and sponsored advertising can be complementary sources of platform profit.

A further concern arises when paid and organic prominence coexist. Organic recommendations provide business users with a free substitute for sponsored visibility. A

platform that earns revenue from paid prominence may therefore have an incentive to reduce the attractiveness of this substitute. It may displace organic results with sponsored positions, reduce the quality of organic recommendations, or otherwise make unpaid comparison less effective. Doing so raises the incremental value of the sponsored position and hence business users' willingness to pay for it. The platform then controls both the quality of unpaid access and the price of the paid alternative.

White (2013) analyzes this trade-off directly. Higher-quality organic results attract users, but they also compete with paid links and can reduce advertising profits. Hence, a profit-maximizing search engine need not choose the organic quality that maximizes user surplus. More recent work sharpens the interaction between the two sides of the recommendation system. Janssen et al. (2026) show that a platform may allocate particularly suitable products to a sponsored position while degrading the organic list, subject to users continuing to participate. The paid position can therefore be relatively informative precisely because the free alternative has been made less informative. Motta and Penta (2026) identify a related displacement effect: purchasing a sponsored position can prevent a rival's organic result from being displayed. In both cases, the appropriate benchmark is not an exogenously given organic ranking. The platform chooses the allocation of paid prominence jointly with the quality or availability of the organic alternative.

This interaction gives rise to a distinct channel of business-user exploitation.²⁰ If organic visibility becomes less effective, sellers may increasingly have to purchase prominence to maintain access to demand. The platform can then appropriate part of the surplus generated by effective access through payments for sponsored visibility. The relevant concern is not that sellers pay for advertising. The platform's design of the unpaid recommendation system can itself increase what business users are willing to pay for the sponsored alternative. Recommendation power allows the platform to affect the scarcity of effective visibility and then monetize that scarcity. Moshary (2025) provides field-experimental evidence on displacement: at a large e-commerce marketplace, exposing users to sponsored listings shifts demand from non-advertised to advertised offers and slightly reduces total purchases, while raising platform earnings per user (see Section 4.4).

The FTC's pending case against Amazon also concerns the degradation of the organic shopping experience. The plaintiffs allege that Amazon increased sponsored placements and displayed less relevant sponsored content, thereby reducing the quality of search results and making advertising increasingly important for sellers seeking to reach customers on Amazon.²¹ These allegations have not been adjudicated. The concern is not the sale of

²⁰ Whether this amounts to an abuse is a further question. In the EU, Article 102(a) TFEU prohibits the imposition of unfair trading conditions on trading partners, providing an established legal basis for addressing exploitative conduct. A theory that a platform degrades unpaid visibility to induce business users to purchase prominence may therefore be assessed within this framework. Where the exercise of market power becomes unfair, however, remains a separate and unsettled question. Not all jurisdictions recognize exploitative abuse as a standalone category; exploitative conduct does not, as such, constitute monopolization under Section 2 of the U.S. Sherman Act.

²¹ FTC et al. v. Amazon.com, Inc., Case No. 2:23-cv-01495-JHC, Second Amended Complaint, ECF No. 327, filed 31.10.2024 (W.D. Wash.).

advertising as such, but that changes in the shopping interface may increase sellers' dependence on paid visibility.

Disclosure addresses a different issue. A “sponsored” label may help users distinguish paid placement from an organic recommendation. It does not, however, restore an organic result that has been displaced, improve an organic ranking that has been degraded, or reduce sellers' dependence on paid prominence. The central question for an exploitative theory is therefore whether sponsored prominence provides additional access to users or whether the platform increases the value of sponsored prominence by worsening the unpaid alternative. In the latter case, recommendation power can be used to extract surplus from business users by making effective access increasingly conditional on payment.

3.2.3 Recommendation as an instrument of consumer-surplus extraction

A third exploitative mechanism arises when recommendation design is used to extract surplus from consumers rather than to favor particular sellers or to monetize prominence. Before purchase, consumers do not observe their match value, whereas the recommender has information about that value and controls what is conveyed through the recommendation. This informational advantage creates scope for a profit-maximizing recommender to distort the information provided to consumers (Peitz and Sobolev, 2025, 2026a, 2026b).²²

Consider a seller that charges a uniform price to consumers who differ in their match value. A fully informative recommendation would advise purchase when the consumer's valuation exceeds the price. A profit-maximizing recommendation need not be fully informative. By pooling some relatively poor matches with good matches, the seller can recommend purchase more often while preserving consumers' willingness to follow a recommendation. Some consumers may consequently buy even though their valuation is below the price. This does not rely on consumer mistakes: given the information conveyed by the recommendation, following it is individually rational (Peitz and Sobolev, 2026a, 2026b).

Price and recommendation policy are jointly determined. More inflated recommendations expand the set of consumers induced to buy but make a positive recommendation less informative. To maintain consumers' willingness to follow it, the seller may then have to lower its price. Conversely, a more selective recommendation can support a higher price. Consumer heterogeneity is important for this mechanism. Consumers may differ in the distributions from which their match values are drawn (Peitz and Sobolev, 2026a, 2026b), or in the information they possess about those match values before purchase (Peitz and Sobolev, 2025). With a common price across such heterogeneous consumers, recommendation design can partly substitute for price discrimination: the seller extracts surplus not only through the price but also through the information conveyed to consumers. The constraint on price discrimination is important. If prices can be conditioned sufficiently finely on the information

²² The mechanism relates to the broader literature on information design and Bayesian persuasion; see Kamenica and Gentzkow (2011) and Rayo and Segal (2010). Roesler and Szentes (2017) and Ennuschat (2025) study the interaction between information structures and monopoly pricing.

available to the seller, the incentive to distort recommendations may disappear because surplus can instead be extracted through prices.

Consumer-surplus extraction and allocative inefficiency should be distinguished. A biased recommendation can reduce consumer surplus without reducing total surplus, while in other circumstances it can induce inefficient trades or prevent efficient ones. Any welfare assessment must consequently account for both the recommendation policy and the induced price response.

This mechanism should be interpreted cautiously as a theory of abuse. The use of information to appropriate consumer surplus does not by itself establish an antitrust concern, just as price discrimination is not generally regarded as abusive merely because it transfers surplus from consumers to the firm.²³ Competition law does recognize direct forms of consumer exploitation, most prominently in the excessive-pricing literature, but there too the difficulty lies in identifying when ordinary exercise of market power crosses the line into abuse.²⁴ The analysis here therefore identifies a potential exploitative use of recommendation power, but does not by itself provide a criterion for when such exploitation constitutes an abuse. Drawing that line would require an additional standard analogous in spirit to that applied in excessive-pricing cases. Whether the recommendation policy generates offsetting benefits and whether possible interventions improve on the outcome are separate questions.

The theories of harm considered in this section identify different ways in which the use of recommendation power may harm competition or consumers. It may be used to weaken competition among offers, to protect the platform's position as a gateway, or to extract surplus from end users or business users. These mechanisms do not depend on self-preferencing and need not involve the same recommendation practices.

4. Testing theories of harm from recommendation power

Section 3 presented specific theories of harm through which recommendation power may lead to foreclosure or exploitation. This section asks how these theories can be assessed empirically. Differences in ranking and prominence are observable, but on their own they reveal neither the mechanism that produces them nor the harm at issue. Four questions therefore have to be separated: how recommendation and the visibility it creates are measured, and how far that visibility affects user choice (Section 4.1); the counterfactual against which observed recommendation is compared, and the benchmark that counterfactual embodies (Section 4.2); why recommendation treatment differs across offers (Section 4.3); and whether that treatment produces the harm posited by the theory of harm (Section 4.4).

²³ On personalized pricing as a possible exploitative abuse and the limited enforcement practice, see Botta and Wiedemann (2020); see also OECD (2018). For a recent argument that AI-enabled personalized price discrimination by dominant firms may constitute exploitative abuse under Article 102(a) TFEU, see Marinova and Bergqvist (2026).

²⁴ On excessive pricing as an exploitative abuse and the difficulties of identifying an administrable intervention threshold, see Evans and Padilla (2005) and Motta and de Streel (2006).

4.1 Measuring recommendation and its importance for user choice

The first step is to measure the recommendation treatment itself. Ordinal rank is often convenient, but it is only one dimension of visibility. Recommendation also operates through the probability that an offer appears at all, its position on the screen, inclusion in a recommendation module, eligibility for a Buy Box or similar featured offer, a badge or rating signal, and the amount of screen space devoted to it. These dimensions need not move together, and none of them maps into exposure in a stable way. A change from rank two to rank one may have a large effect on exposure, whereas a one-position change far down a list may have almost none, and the same numerical rank may have different consequences on mobile and desktop interfaces or when one result is accompanied by images, labels, or other salient information. Jürgensmeier and Skiera (2025), for example, construct a platform-wide measure of nonpersonalized product visibility across search terms rather than equating visibility with a product's position in a particular search.

The economically relevant object is therefore not rank as such but the visibility that recommendation design creates, in the broad sense of Section 2.2: how likely users are to consider an offer and how they evaluate it. This is the measurement counterpart of the distinction between formal and effective access developed in Section 2.2: a seller can remain listed while receiving little of the demand that the platform mediates.

How much that exposure matters for user choice is a separate question. Evidence on recommendation should therefore be combined, where possible, with evidence on clicks, traffic, conversion, or transactions. Aguiar and Waldfogel (2021) estimate the effect of inclusion in Spotify's algorithmic Top 50 playlists on streams, using a regression-discontinuity design that compares songs just inside and just outside the threshold. They also employ an instrumental-variables strategy for the curated New Music Friday playlist and find large effects on song success. Using experimental variation in hotel rankings at Expedia, Ursu (2018) finds that rankings affect which offers consumers search but not their purchasing decisions conditional on search. Differences in visibility matter economically only to the extent that prominence changes user behavior.

4.2 Constructing the counterfactual

Any assessment of recommendation requires a counterfactual against which the observed treatment is compared. Two counterfactuals should be distinguished. The first concerns recommendation treatment: holding the observed market environment fixed, how would the same offer have been displayed absent the characteristic suspected of affecting recommendation? The second concerns the market outcome under an alternative recommendation rule. Prices, quality, entry, participation, and investment need not remain unchanged when visibility rules change. As Section 2.1 showed, a rule that matches users well for a given set of offers need not remain attractive once sellers respond to it, so the two counterfactuals can point in different directions.

The two questions are also addressed at different points in the analysis. Whether recommendation responds to the suspected characteristic is a question about treatment, taken

up in Section 4.3. What would follow from a different rule is a question about welfare, taken up in Section 4.4. Both, however, presuppose a benchmark, and the choice of benchmark is not given by the data.

There is no uniquely neutral ranking against which every observed recommendation can be judged. Empirical approaches therefore embody a choice of benchmark. Aguiar, Waldfogel and Waldfogel (2021) use an outcome-based test: conditional on Spotify's ex ante assessment of songs, differences in their subsequent success can reveal whether particular types of music were systematically favored or disfavored. The benchmark is the platform's own assessment: songs it judged equally promising should, absent favoritism, go on to perform equally well, so systematic differences in outcomes among equally ranked songs indicate that placement responded to something other than expected success. Reimers and Waldfogel (2023) instead define bias relative to a frontier that maximizes a weighted sum of consumer and seller surplus. As Section 2.1 noted, varying that weight traces a one-parameter family of rankings, corresponding to different values of α in the scoring rule. Their definition of bias does not select among them: a ranking is unbiased if it lies anywhere on the frontier, and biased if it delivers an outcome interior to it. In their specification, bias enters as an additional term in the ranking index attached to platform-favored products, so that departures from the frontier can be estimated without committing to a view about the appropriate weight on consumer and seller surplus. What drives differential treatment, and therefore what the analysis must isolate, depends on the theory being tested, as the next subsection makes concrete.

4.3 Distinguishing the suspected mechanism from benign explanations

The next step is to establish why the recommendation treatment changes. A correlation between prominence and a platform-relevant characteristic is not sufficient, because such characteristics are usually correlated with features of an offer that bear on user value: own offers may differ systematically in quality or availability, sellers using platform logistics may deliver faster, and offers carrying high commissions may be more popular with users. Each of these may provide a benign explanation for the difference in treatment. The empirical question is whether, after accounting for these benign explanations, visibility changes because of the characteristic underlying the theory of harm.

The relevant characteristic depends on the theory. For self-preferencing it may be ownership or seller identity. For profit-driven steering it may be the platform's return from the transaction or a sponsored payment. For gateway protection it may be an off-platform price, use of a direct channel, exclusivity, or participation in a competing service. Where recommendation is tied to the use of another platform service, it may be the adoption of the platform's logistics, payment, or other ancillary service. Credible variation in these characteristics can come from stockouts, changes in platform policy or algorithm design, eligibility thresholds, regulatory interventions, discontinuities, or variation within the same product or seller over time.

Empirical studies of Amazon illustrate different ways of relating prominence to ownership or seller identity. Chen and Tsai (2024) exploit Amazon stockouts to compare recommendations for the same product when Amazon is and is not a seller, generating within-product variation

in seller identity. Farronato, Fradkin and MacKay (2023) compare the search positions of Amazon-branded and other products conditional on observed product characteristics. Waldfogel (2024) uses more than eight million search results across 22 Amazon domains and documents a substantial own-product ranking advantage, together with a decline in that advantage following Amazon's designation as a gatekeeper under the DMA while the rankings of other major brands did not change comparably. Dendorfer and Seibel (2026) compare the same products on Amazon's US and Canadian domains: the "Similar item to consider" box shows products of other brands in the US but almost exclusively Amazon Basics products in Canada, a difference not explained by observed product characteristics or sales. They estimate that products denied a recommendation in Canada sell at least 12% less.

For gateway-protection theories, the characteristic associated with prominence may instead concern conduct on competing channels. The evidence discussed in Section 3.1.2 illustrates the approach: Hunold, Kesler and Laitenberger (2020) relate a hotel's position in an online travel agent's search results to the prices the hotel charges on competing channels. The link may also be indirect: as Bar-Isaac and Shelegia (2026) show in their model, off-platform pricing can affect prominence through its effect on on-platform sales, and the same applies to conversion rates (Section 3.1.2), so a test that only checks whether off-platform prices enter the ranking rule may miss the channel through which off-platform conduct is disciplined.²⁵ Cure et al. (2022) study the vertically integrated meta-search platform Kayak and find that, for a given hotel, offers from affiliated OTAs are more visible than equally priced offers from unaffiliated OTAs; they also find hotels to be less prominent when the unaffiliated OTA Expedia offers the lowest price.

The benign explanations noted above are not the only ones. Recommendation may also respond to capacity constraints, fraud risk, safety, or experimentation and learning. As discussed in Section 2.1, even a user-welfare-maximizing recommender may depart from the best current match to induce experimentation (Che and Hörner, 2018). These explanations are harder to handle than differences in price or quality, because they are often unobserved by outsiders and may vary over time.

The Kakao Mobility case illustrates this difficulty in an enforcement setting. The KFTC found that Kakao's dispatch algorithm favored affiliated franchise taxis and treated this as part of a strategy to strengthen Kakao's position in the taxi-franchise market. The Seoul High Court overturned the decision, finding that unlawful discrimination had not been established and

²⁵ Changes that Booking.com made after a regulatory dialogue with the European Commission under the Digital Markets Act illustrate this point. Booking.com stated that it does not use prices outside its platform in its default ranking or for eligibility for its visibility-enhancing partner programs (Genius, Preferred, and Preferred Plus) and, since September 2026, no longer uses them for eligibility for its Booking Sponsored Benefit program; previously, external prices had been taken into account for eligibility for this program, under which Booking.com lowers a listing's displayed price at its own expense (European Commission, Factsheet "How the DMA ensures businesses using Booking.com are free to set their prices on and off Booking.com", 28 September 2026). This removes the direct link between off-platform prices and visibility but not, by itself, the indirect one: to the extent that the ranking rewards on-platform sales or conversion, lower prices elsewhere can still reduce a property's visibility. The evidence of Hunold, Kesler and Laitenberger (2020), which includes Booking.com, refers to an earlier period.

that the authority had not sufficiently demonstrated anticompetitive effects.²⁶ The dispute illustrates the difficulty of distinguishing strategically motivated preferential steering from operational determinants of allocation when both can produce systematic differences in treatment.

4.4 Establishing the harm

Establishing that recommendation responds to the characteristic posited by the theory of harm does not complete the analysis. The final step is to establish the harm that the theory posits. The required evidence differs across the mechanisms in Section 3.

For foreclosure through competition among offers, diversion of traffic or sales provides evidence of the first link in the mechanism but is not sufficient. The question is whether the disadvantaged suppliers impose a competitive constraint and whether the recommendation practice weakens that constraint. Evidence may therefore concern effects on prices, quality, seller participation, investment, or entry as well as sales. Because foreclosure can be partial, the fact that rivals remain on the platform does not resolve the issue; what matters is whether their reduced effective access weakens the competitive constraints faced by the favored offer.

The Korean Naver Shopping case illustrates how demanding this assessment can be in practice.²⁷ Naver's comparison-shopping service had favored merchants on its own marketplace over those on rival marketplaces, and Naver's marketplace gained share and turnover relative to its rivals. The Supreme Court of Korea nevertheless held that anticompetitive effects had not been sufficiently established and overturned the judgment upholding the KFTC's finding of abuse. In the court's assessment, rival marketplaces obtained most of their transactions through direct visits rather than through Naver Shopping, their transaction volumes continued to grow, it had not been shown that their costs increased, and entry continued, leaving it unclear whether Naver's gains reflected the challenged conduct rather than competition on the merits or overall market growth.

For gateway protection, the relevant effect is on the viability of an alternative route to users. Immediate traffic diversion may be informative, but the theory often turns on subsequent changes in off-platform transactions, multi-homing, direct-channel investment, participation in a rival gateway, or the scale available to a complementary service such as logistics. These effects can unfold over time. Reduced traffic today may mean less data, weaker learning, lower reputation, or less participation tomorrow, which in turn can reduce the alternative route's ability to constrain the platform. Where recommendations depend on past user interaction, lower current visibility may also translate into lower future visibility because it generates fewer interactions on which subsequent recommendations are based (Fletcher,

²⁶ Korea Fair Trade Commission, 13.6.2023, Decision No. 2023-093, Case No. 2021Sigam0760, Kakao Mobility; Seoul High Court, 22.5.2025, Case No. 2023Nu50891, Kakao Mobility v Korea Fair Trade Commission. An appeal is pending before the Supreme Court of Korea, Case No. 2025Du34036. See Franck and Peitz (2026).

²⁷ Korea Fair Trade Commission, 27.1.2021, Decision No. 2021-027, *Naver Shopping*; Seoul High Court, 14.12.2022, Case No. 2021Nu36129; Supreme Court of Korea, 16.10.2025, Case No. 2023Du32709.

Ormosi and Savani, 2023). Establishing such dynamic effects empirically is more demanding than measuring short-run diversion.

For exploitative uses, a weakening of competitive constraints need not be shown. The relevant question is whether recommendation allows the platform to obtain additional surplus at the expense of end users or business users. Depending on the theory, evidence may concern whether users are steered toward offers that yield the platform a higher return but provide lower expected user value; whether sellers make additional payments for prominence because organic visibility has become less effective; or whether information design changes users' choices in a way that transfers surplus to the platform.

For each of these mechanisms, the welfare consequences of a different recommendation rule depend on how market participants respond to it. This is the second counterfactual of Section 4.2, and Section 2.1 has already set out why it matters: as the analyses of Dinerstein et al. (2018) and Rhodes and Zhou (2026) discussed in Section 2.1, show, recommendation design affects sellers' pricing incentives as well as the quality of the match. The point applies to the welfare frontier discussed in Section 4.2 as well. Reimers and Waldfogel (2023) hold supplier prices fixed when they compute debiased rankings, and note that suppliers might reprice if they faced a different ranking rule. Because the frontier is traced by rankings that weight the price-cost margin against rank-independent mean utility, and the latter itself depends on price,²⁸ the frontier is defined relative to the prices that the observed rule induces, and a modified ranking rule need not leave it where it was.

Moshary (2025) provides evidence on the displacement of organic by sponsored listings (Section 3.2.2) from a field experiment at a large e-commerce marketplace, in which sponsored listings were suppressed for a randomly selected 3% of site traffic. Users in that group therefore saw a purely organically ranked page, unlike in experiments that manipulate a single seller's advertising, where rivals' sponsored listings remain visible to users in both arms. Because only a small share of buyers is affected, sellers set their prices, listing decisions, and advertising campaigns with reference to the prevailing environment, in which sponsored search is present. Relative to the organically ranked group, exposure to sponsored listings raises spending on advertised offers by about 17% but reduces spending on non-advertised offers by about 8%. Aggregated across listings, the loss dominates: total spending on the platform falls by roughly 1.7% and the probability of any purchase by about 0.5%, so sponsored prominence reallocates demand rather than expanding it. Platform earnings per user are nonetheless about 4% higher, because advertising revenue more than offsets the forgone commissions. Monetizing prominence can therefore be privately profitable even where it reduces the volume of matches the platform intermediates.²⁹

²⁸ Reimers and Waldfogel (2023) assume zero marginal costs, so that price equals per-unit variable profit, and note that with observed costs one would substitute per-unit profits throughout. This is the assumption under which the scoring coefficient α of Section 2.1 maps into the weight on consumer and seller surplus. Where marginal costs differ across offers and are unobserved, seller surplus cannot be recovered from price and the frontier is not identified.

²⁹ A decline in transactions need not indicate worse matching if sponsored prominence steers users away from purchases they would regret. Moshary (2025) finds mixed evidence: users exposed to sponsored listings complain slightly less per transaction, which points to better matches, but are less likely to purchase a second

Randomization is at the buyer level, so each seller posts one price and makes one listing decision that applies to both groups. The estimates are therefore obtained holding sellers' decisions fixed, as in the supply-taken-as-given assessment of Section 2.1. Moshary's back-of-the-envelope calculation shows how a price response to a platform-wide ban would change the estimates. If prices are 5% higher with sponsored search than without and platform-level demand is inelastic (an elasticity of -0.5), the implied fall in purchase probability rises from about 0.5% to about 3%. The effect on spending, by contrast, turns from -1.7% to $+0.6\%$, because the higher price more than offsets the fall in purchases.³⁰

Two structural studies of Amazon assess the welfare implications while allowing sellers to respond. Lee and Musolff (2026) estimate a model of the Buy Box that jointly accounts for consumer consideration, algorithmic prominence, seller pricing, and entry. They estimate that 95% of consumers consider only the Buy Box offer and find that the algorithm favors offers sold by Amazon: such offers are featured as frequently as observably similar third-party offers priced about 5% lower. Yet consumers also prefer Amazon's offers, consumer surplus is approximately maximized at the estimated degree of Amazon promotion, and their long-run counterfactuals, in which seller prices and entry adjust to the recommendation rule, provide no evidence for the barrier-to-entry theory of harm. Lam (2022) obtains a related result from a structural model of search and pricing on Amazon: replacing the observed product arrangement with a more neutral one shifts profits from Amazon to third-party sellers but reduces consumer welfare as consumers adjust their search and sellers adjust prices. Preferential prominence, even where convincingly established, therefore does not determine the welfare effects of removing it.

There is consequently no single empirical test for whether the exercise of recommendation power harms competition or consumers. The evidence has to be organized around the particular causal chain proposed by the theory of harm. This discipline avoids treating every observed difference in visibility as harmful, but it also avoids reducing the inquiry to whether rivals remain formally available to users. Even a complete account of the causal chain establishes only who gains and who loses from a recommendation practice; whether the outcome should be considered harmful is a separate question. The same practice may lower prices for end users while restricting business users' pricing, or affect consumers with different information or preferences in different ways (Section 3.2.3). Whether losses to some users should be offset by gains to others is a normative question governed by the applicable legal standard, and existing standards answer it differently.

A standard defined in terms of surplus must specify how the welfare of different user groups is weighted; a standard focused on protecting the competitive process need not require a

time, which does not. Adjusting for returns, on the assumption that they fall in proportion to complaints from a base rate of 18%, reduces the estimated decline in purchase probability from 0.47% to 0.32%, so fewer regretted purchases explain only part of the decline.

³⁰ Prices may be higher with sponsored search because ad fees raise sellers' marginal cost of a sale and because sponsored search may soften price sensitivity. Evidence for the latter is suggestive: users exposed to sponsored listings buy more expensive sponsored items and users exposed to sponsored listings buy more expensive sponsored items and are less likely to reorder search results so that the cheapest items appear first.

comprehensive balancing of gains and losses, which makes it more readily applicable to foreclosure theories than to exploitation, where the distribution of surplus is itself central.³¹

5. Discussion and conclusion

Recommendation systems create value by selecting, ordering, and presenting offers: they can improve matching, induce lower prices or higher quality, and facilitate learning. The same control over visibility becomes recommendation power where users rely on a platform for discovery and evaluation and business users depend on it to reach those users, because differences in visibility then become differences in effective access. Section 3 showed that recommendation power can be used to weaken competition among offers, to protect the platform's position as a gateway, or to extract surplus from end users or business users. What distinguishes these theories is not the recommendation practice but the mechanism through which it affects competition or surplus, and the favored offer need not be the platform's own. Section 4 showed that an observed visibility advantage must be linked both to the characteristic posited by the theory of harm, rather than to differences in price, quality, or other characteristics relevant to users, and to the competitive or exploitative effect at issue. Even where preferential treatment is clearly established, removing it need not benefit consumers, as sellers adjust prices, quality, and entry to the changed recommendation rule.

For competition policy, the first implication concerns the meaning of access. Formal access to a platform can coexist with very limited access to demand. A seller that remains listed but rarely appears in recommendations, loses eligibility for a featured offer, or is placed where users seldom look may retain contractual access while losing much of the commercial value of participation. This matters particularly where users concentrate their search on the platform

³¹ How countervailing effects on different user groups are weighed depends on the jurisdiction. US case law has rejected the justification of anticompetitive effects in one market by procompetitive effects in another, notably in merger control (*United States v. Philadelphia National Bank*, 374 U.S. 321, 370 (1963)). This makes market definition particularly important: by defining a single market for credit-card transactions, the US Supreme Court in *Ohio v. American Express Co.*, 585 U.S. 529 (2018), required the plaintiffs to establish anticompetitive effects in the market as a whole, taking into account the effects on both merchants and cardholders. Under Article 101(1) TFEU, the European Court of Justice has held that market definition must not prevent the consideration of interactions between the two sides of a payment card system when assessing whether an agreement restricts competition by object (Case C-67/13 P, *Groupeement des cartes bancaires v Commission*, ECLI:EU:C:2014:2204, paras. 76–79). Moreover, the Court of Justice has recognized that efficiencies on both sides of a two-sided payment card system may be taken into account, provided that appreciable objective advantages also arise for the consumer group affected by the restriction, even if those advantages are insufficient on their own to offset the restrictive effects (Case C-382/12 P, *MasterCard v Commission*, ECLI:EU:C:2014:2201, paras. 237 and 241). The European Commission generally limits the balancing of anticompetitive effects against efficiencies to the relevant market in which the restriction occurs, allowing efficiencies arising in related markets to be taken into account only exceptionally, where the consumers affected by the restriction and those benefiting from the efficiencies are substantially the same. See *Guidelines on Horizontal Co-operation Agreements*, OJ C 259, 21.7.2023, 1, paras. 583–584. See Franck and Peitz (2021, pp. 125–126) who argue that, in the case of two-sided platforms, whether effects on different user groups may be balanced against each other should not depend on whether the platform's two sides are assigned to a single relevant market or to separate markets. The substantive conditions for such balancing should be distinguished from the question of market definition. South African competition law illustrates the additional complexity where competition policy explicitly pursues multiple objectives, including ensuring that small and medium-sized enterprises have “an equitable opportunity to participate in the economy” (Competition Act 89 of 1998, s. 2(e)). Assessing the effects of a practice may then require balancing competing objectives that are not readily commensurable.

and business users have weak alternatives for reaching them. As argued in Section 2.2, recommendation power is not a separate form of market power; assessing it requires evidence both on how far users rely on the platform's recommendations and on how far business users depend on the platform for access to demand.

A second implication concerns the benchmark against which recommendation practices are assessed. There is no general economic presumption in favor of equal exposure or a single "neutral" ranking. A rule based on expected match value, consumer surplus, price, quality, or some combination of these criteria can produce different rankings, and a rule that is optimal taking existing offers as given need not remain optimal once sellers respond to it. This makes the theory of harm central both to determining which recommendation practices should be prohibited and to designing adequate remedies. If the concern is a visibility penalty linked to lower off-platform prices, the relevant intervention is directed at that conditioning of visibility, not at equalizing the ranks of all sellers. If the concern is preferential treatment of an own offer, removing the ownership-related advantage need not require eliminating legitimate differences in price, quality, or service performance. More generally, a remedy should target the mechanism that produces the identified harm rather than replace the platform's recommendation rule with an administratively specified ranking.

This point also limits the usefulness of broad nondiscrimination requirements. Equal treatment can remove harmful favoritism, but it can also prevent recommendation from responding to economically relevant differences among offers. Conversely, transparency alone may reveal that a ranking criterion exists without changing the incentives created by it. The appropriate remedy therefore depends on why the recommendation rule is harmful. Changes in ranking criteria, eligibility conditions, sponsored placement, or the use of ancillary services can alter seller pricing, quality, participation, and investment. Remedy analysis should account for these responses just as the liability analysis does. The same lesson emerges from the literature on search design: recommendation is not merely an allocation of a fixed amount of visibility among fixed offers; it can change the behavior that determines which offers are available and on what terms.

Recommendation-based conduct also creates practical difficulties for ex post enforcement. Ranking and recommendation rules can be modified continuously, and treatment may differ across users, queries, devices, and time. The platform generally observes the inputs to the recommendation algorithm, its internal performance measures, and the experiments used to modify the system, whereas business users and authorities observe only part of the resulting treatment. Section 4 showed why this informational asymmetry matters. Establishing whether seller identity, commissions, off-platform pricing, sponsored payments, or another platform-relevant variable affects visibility often requires detailed data or credible changes in those variables. Establishing dynamic harm is harder still. Where reduced traffic affects data, reputation, participation, or the viability of an alternative gateway, short-run diversion provides only part of the relevant evidence. Recommender systems can add another feedback channel where current visibility affects user interaction and, through it, the information used for future recommendations.

These difficulties provide a rationale for rules that reduce the evidentiary burden where particular risks are considered sufficiently important. The prohibition of self-preferencing in ranking under Article 6(5) of the Digital Markets Act (DMA) is one such response (Peitz, 2023). It prohibits a particular form of self-preferencing by designated gatekeepers without requiring an individualized welfare analysis. This improves administrability, but it also excludes an equilibrium-based justification under which the ranking rule induces seller responses that ultimately benefit consumers (Franck and Peitz, 2026). The trade-off illustrates why a categorical rule and the effects-based analysis developed in this paper need not lead to the same assessment. The theories considered here are broader than the conduct covered by Article 6(5) of the DMA: they include preferential treatment among third parties, visibility penalties linked to conduct on competing channels, and exploitative recommendation practices. Some of these practices may nevertheless be constrained by other DMA provisions, which concern, among other things, data inputs into recommendation, choice architecture and defaults, conditions of access, and the use of recommendation design to circumvent substantive obligations (Franck and Peitz, 2026). The DMA thus addresses recommendation power selectively through a set of specific obligations rather than through a general rule governing recommendation design.

There is ample room for further economic research. First, empirical work is considerably more developed for own-product advantages in rankings than for recommendation used to protect an existing gateway. The latter theories require evidence linking visibility to conduct on competing channels and then linking that treatment to investment in direct distribution, multi-homing, participation in rival intermediaries, or the scale of competing services. The hotel-booking evidence on ranking and off-platform pricing provides one example, but comparable evidence remains limited in other platform settings. A particularly useful empirical design would exploit changes in recommendation treatment triggered by off-platform prices, exclusivity, or the adoption of platform services and trace their effects beyond immediate platform sales.

Second, more theoretical and empirical work is needed on the longer-run effects of recommendation systems on business-user behavior and the evolution of competing routes to users. Existing work shows that seller pricing, entry, and participation respond to search design and recommendation rules. Much less is known about responses involving quality investment, multi-homing, direct-channel investment, and the accumulation of reputation or data. These responses are especially important for gateway-protection theories, where a current loss of traffic may affect the future ability of an alternative route to compete.

Third, there is relatively little evidence on the effects of alternative remedies once a harmful recommendation practice has been identified. Removing the feature that gives rise to the harm need not have the same effects as imposing a broader change in ranking or prominence. Sellers and users may respond differently across interventions, and these responses can affect prices, quality, participation, and platform incentives. Regulatory interventions and platform policy changes can provide useful variation for comparing such responses and for identifying remedies that address the harmful mechanism without eliminating beneficial properties of recommendation.

Recommendation systems are valuable precisely because they do not treat all available options alike. That feature is also what makes recommendation power economically important. The relevant competition question is whether the platform uses its control over recommendation to weaken a competitive constraint or to exploit end users or business users. Answering that question requires identifying the mechanism and the economically relevant counterfactual. Keeping the mechanism at the center of the analysis provides a way to assess recommendation-based conduct without treating recommendation systems per se as suspect.

References

- Aguiar, L., Waldfogel, J., 2021. Platforms, power, and promotion: Evidence from Spotify playlists. *Journal of Industrial Economics* 69, 653–691.
- Aguiar, L., Waldfogel, J., Waldfogel, S., 2021. Playlisting favorites: Measuring platform bias in the music industry. *International Journal of Industrial Organization* 78, 102765.
- Armstrong, M., Zhou, J., 2011. Paying for prominence. *Economic Journal* 121, F368–F395.
- Athey, S., Ellison, G., 2011. Position auctions with consumer search. *Quarterly Journal of Economics* 126, 1213–1270.
- Athey, S., Scott Morton, F., 2022. Platform annexation. *Antitrust Law Journal* 84, 677–703.
- Bakos, J.Y., 1997. Reducing buyer search costs: implications for electronic marketplaces. *Management Science* 43, 1676–1692.
- Balasingham, B., D’Amico, A.S., 2025. Beyond exclusion: revisiting exploitative abuses in digital platform markets. *Journal of Antitrust Enforcement* 13, 265–283.
- Bar-Isaac, H., Shelegia, S., 2025. Monetizing steering. *Journal of Political Economy Microeconomics*, in press. <https://doi.org/10.1086/740100>.
- Bar-Isaac, H., Shelegia, S., 2026. Impressing the algorithm: Sales-based ranking, learning, and off-platform pricing. Unpublished manuscript.
- Belleflamme, P., Peitz, M., 2019. Managing competition on a two-sided platform. *Journal of Economics and Management Strategy* 28, 5–22.
- Belleflamme, P., Peitz, M., 2021. *The Economics of Platforms: Concepts and Strategy*. Cambridge University Press.
- Botta, M., Wiedemann, K., 2020. To discriminate or not to discriminate? Personalised pricing in online markets as exploitative abuse of dominance. *European Journal of Law and Economics* 50, 381–404.
- Bougette, P., Budzinski, O., Marty, F., 2022. Self-preferencing and competitive damages: a focus on exploitative abuses. *Antitrust Bulletin* 67, 190–207.

- Bourreau, M., Gaudin, G., 2022. Streaming platform and strategic recommendation bias. *Journal of Economics & Management Strategy* 31, 25–47.
- Burguet, R., Caminal, R., Ellman, M., 2015. In Google we trust? *International Journal of Industrial Organization* 39, 44–55.
- Calvano, E., Calzolari, G., Denicolò, V., Pastorello, S., 2023. Artificial intelligence, algorithmic recommendations and competition. CEPR Discussion Paper No. 18176, Centre for Economic Policy Research.
- Carlton, D., Waldman, M., 2002. The strategic use of tying to preserve and create market power in evolving industries. *RAND Journal of Economics* 33, 194–220.
- Carroni, E., Madio, L., Shekhar, S., 2024. Superstar exclusivity in two-sided markets. *Management Science* 70, 991–1011.
- Che, Y.-K., Hörner, J., 2018. Recommender systems as mechanisms for social learning. *Quarterly Journal of Economics* 133, 871–925.
- Chen, N., Tsai, H.-T., 2024. Steering via algorithmic recommendations. *RAND Journal of Economics* 55, 501–518.
- Chen, Y., He, C., 2011. Paid placement: advertising and search on the Internet. *Economic Journal* 121, F309–F328.
- Choi, J.P., Jeon, D.-S., Whinston, M.D., 2026. Tying with network effects. *American Economic Review* 116, 332–374.
- Choi, J.P., Peitz, M., 2018. You are judged by the company you keep: reputation leverage in vertically related markets. *International Journal of Industrial Organization* 61, 351–379.
- Crémer, J., de Montjoye, Y.-A., Schweitzer, H., 2019. *Competition Policy for the Digital Era*. Report for the European Commission.
- Cure, M., Hunold, M., Kesler, R., Laitenberger, U., Larrieu, T., 2022. Vertical integration of platforms and product prominence. *Quantitative Marketing and Economics* 20, 353–395.
- de Cornière, A., Taylor, G., 2014. Integration and search engine bias. *RAND Journal of Economics* 45, 576–597.
- de Cornière, A., Taylor, G., 2019. A model of biased intermediation. *RAND Journal of Economics* 50, 854–882.
- Dendorfer, F., Seibel, R., 2026. What’s in the box? The effect of self-preferencing on Amazon Marketplace. Unpublished manuscript.
- Dinerstein, M., Einav, L., Levin, J., Sundaresan, N., 2018. Consumer price search and platform design in internet commerce. *American Economic Review* 108, 1820–1859.
- Eliaz, K., Spiegel, R., 2011. A simple model of search engine pricing. *Economic Journal* 121, F329–F339.

- Ennuschat, P., 2025. Targeted information design: Shaping monopolistic markets with heterogeneous consumers. Unpublished manuscript, Universitat Pompeu Fabra.
- Etro, F., 2023. Hybrid marketplaces with free entry of sellers. *Review of Industrial Organization* 62, 119–148.
- Etro, F., 2024. E-commerce platforms and self-preferencing. *Journal of Economic Surveys* 38, 854–878.
- European Commission, 2023. Out-of-market efficiencies in competition enforcement – Note by the European Union. OECD Document DAF/COMP/WD(2023)110.
- Evans, D.S., Padilla, A.J., 2005. Excessive prices: using economics to define administrable legal rules. *Journal of Competition Law & Economics* 1, 97–122.
- Farronato, C., Fradkin, A., MacKay, A., 2023. Self-preferencing at Amazon: evidence from search rankings. *AEA Papers and Proceedings* 113, 239–243.
- Fletcher, A., Ormosi, P.L., Savani, R., 2023. Recommender systems and supplier competition on platforms. *Journal of Competition Law & Economics* 19, 397–426.
- Franck, J.-U., Peitz, M., 2021. Market definition in the platform economy. *Cambridge Yearbook of European Legal Studies* 23, 91–127.
- Franck, J.-U., Peitz, M., 2026. Recommendation power: platform objectives and the Digital Markets Act. CRC TR 224 Discussion Paper No. 781.
- Fumagalli, C., Motta, M., 2020. Dynamic vertical foreclosure. *Journal of Law and Economics* 63, 763–812.
- Hagiü, A., Jullien, B., 2011. Why do intermediaries divert search? *RAND Journal of Economics* 42, 337–362.
- Hagiü, A., Teh, T.-H., Wright, J., 2022. Should platforms be allowed to sell on their own marketplaces? *RAND Journal of Economics* 53, 297–327.
- Hagiü, A., Wright, J., 2023. Data-enabled learning, network effects, and competitive advantage. *RAND Journal of Economics* 54, 638–667.
- Hagiü, A., Wright, J., 2024. Marketplace leakage. *Management Science* 70, 1529–1553.
- Heidhues, P., Köster, M., Köszegi, B., 2023. Steering fallible consumers. *Economic Journal* 133, 1430–1465.
- Hunold, M., Kesler, R., Laitenberger, U., 2020. Rankings of online travel agents, channel pricing, and consumer protection. *Marketing Science* 39, 92–116.
- Janssen, M.C.W., Jungbauer, T., Preuss, M., Williams, C., 2026. Search platforms: big data and sponsored positions. *Economic Journal*, in press. <https://doi.org/10.1093/ej/ueag002>.
- Johnson, J.P., Rhodes, A., Wildenbeest, M., 2023. Platform design when sellers use pricing algorithms. *Econometrica* 91, 1841–1879.

- Jürgensmeier, L., Skiera, B., 2025. Measuring self-preferencing on digital platforms. *Journal of Marketing*, in press. <https://doi.org/10.1177/00222429251360772>.
- Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. *American Economic Review* 101, 2590–2615.
- Kittaka, Y., Sato, S., Zenryo, Y., 2023. Self-preferencing by platforms: A literature review. *Japan and the World Economy* 66, 101191.
- Lam, H.T., 2022. Platform search design and market power. Working paper, UCLA Anderson School of Management.
- Lee, K.H., Musolff, L., 2026. Two-sided markets shaped by platform-guided search. *Econometrica*, in press.
- Long, F., Amaldoss, W., 2024. Self-preferencing in e-commerce marketplaces: The role of sponsored advertising and private labels. *Marketing Science* 43, 925–952.
- Marinova, M., Bergqvist, C., 2026. AI-enabled price discrimination as an exploitative abuse of dominance under EU competition law. *Journal of Competition Law & Economics*, nhag006. <https://doi.org/10.1093/joclec/nhag006>.
- Marsteller, L., 2026. Recommendation design, pricing, and regulation. ECONtribute Discussion Paper No. 430, University of Bonn and University of Cologne.
- Moshary, S., 2025. Does sponsored search advertising augment organic search? Evidence from an e-commerce platform. *Management Science* 71, 9687–9709.
- Motta, M., 2023. Self-preferencing and foreclosure in digital markets: theories of harm for abuse cases. *International Journal of Industrial Organization* 90, 102974.
- Motta, M., de Streel, A., 2006. Excessive pricing and price squeeze under EU law. In: Ehlermann, C.-D., Atanasiu, I. (Eds.), *European Competition Law Annual 2003: What Is an Abuse of a Dominant Position?* Hart Publishing, Oxford, pp. 91–125.
- Motta, M., Penta, A., 2026. Market effects of brand search advertising. *RAND Journal of Economics* 57 (3), 724–744.
- OECD, 2018. *Personalised Pricing in the Digital Era*. OECD Roundtables on Competition Policy Papers No. 222. OECD Publishing, Paris.
- Peitz, M., 2023. How to apply the self-preferencing prohibition in the DMA. *Journal of European Competition Law & Practice* 14, 310–315.
- Peitz, M., Schweitzer, H., 2016. Suchmaschinen unter der Lupe: Informationsherrschaft und ihre Schranken. *Wirtschaftsdienst* 96, 825–829.
- Peitz, M., Sobolev, A., 2025. Biased recommendations and differentially informed consumers. *Economic Theory Bulletin* 13, 245–261.

- Peitz, M., Sobolev, A., 2026a. Inflated recommendations. *RAND Journal of Economics* 57, 368–385.
- Peitz, M., Sobolev, A., 2026b. Recommendation power and competition. *International Journal of Industrial Organization* 108, 103310.
- Rayo, L., Segal, I., 2010. Optimal information disclosure. *Journal of Political Economy* 118, 949–987.
- Reimers, I., Waldfogel, J., 2023. A framework for detection, measurement, and welfare analysis of platform bias. NBER Working Paper No. 31766.
- Rhodes, A., Zhou, J., 2026. Personalization and privacy choice. *RAND Journal of Economics*, in press.
- Roesler, A.-K., Szentes, B., 2017. Buyer-optimal learning and monopoly pricing. *American Economic Review* 107, 2072–2080.
- Schweitzer, H., Haucap, J., Kerber, W., Welker, R., 2018. *Modernisierung der Missbrauchsaufsicht für marktmächtige Unternehmen*. Nomos.
- Schweitzer, H., Welker, R., 2019. Competition policy for the digital era. *CPI Antitrust Chronicle* 3 (2), 16–24.
- Ursu, R.M., 2018. The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science* 37, 530–552.
- Waldfogel, J., 2024. Amazon self-preferencing in the shadow of the Digital Markets Act. NBER Working Paper No. 32299.
- White, A., 2013. Search engines: Left side quality versus right side profits. *International Journal of Industrial Organization* 31, 690–701.
- Zenryo, Y., 2022. Platform encroachment and own-content bias. *Journal of Industrial Economics* 70, 684–710.
- Zhou, B., Zou, T., 2023. Competing for recommendations: The strategic impact of personalized product recommendations in online marketplaces. *Marketing Science* 42, 360–376.