Discussion Paper Series – CRC TR 224

Discussion Paper No. 781
Project B 05

Recommendation Power: Platform Objectives and the Digital Markets Act

Jens-Uwe Franck[1]
Martin Peitz[2]

September 2026

[1]Jens-Uwe Franck: University of Mannheim, Department of Law, and MaCCI.
[2]University 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.

Recommendation Power:

Platform Objectives and the Digital Markets Act

Jens-Uwe Franck and Martin Peitz*

September 10, 2026

Abstract

Digital platforms allocate scarce attention through rankings, recommendations, defaults, and prominence. This article develops the concept of recommendation power as the point of intersection between two established concepts: information power and intermediation power, which is the capability to actively steer user choice by a platform that is a bottleneck for effective access. We propose a taxonomy of the objectives that typically guide platforms' recommendation design, distinguishing an independent-advice benchmark from fifteen other objectives that may motivate systematic departures from it, grouped into direct commercial and data objectives, indirect and strategic objectives, and objectives relating to influence, welfare, and integrity. We then analyze how the EU's Digital Markets Act (DMA) constrains recommender design, though it does not regulate recommendation power as such. Through obligations on ranking and data and transparency duties as well as through its anticircumvention provisions, the DMA protects certain user choices and access rights from the exercise of recommendation power.

JEL codes: K21, L15, L40, L51, L86.

Keywords: recommender systems, recommendation bias, ranking, self-preferencing, gatekeepers, Digital Markets Act, platform regulation, steering.

* 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.

TABLE OF CONTENTS

I. INTRODUCTION 4
II. THE CONCEPT OF RECOMMENDATION POWER 6
A. Information Power 6
B. Intermediation Power 6
C. Recommendation Power: the Intersection 6
D. Platforms as Private Rule-Setters and the Limits of Ex Post Control 7
E. Summary 9
III. A TAXONOMY OF RECOMMENDATION OBJECTIVES 9
A. What Is Being Classified—and Relative to Which Benchmark? 9
Category 1. The Independent-Advice Benchmark 10
B. Objectives Aimed at Immediate Value Capture (Categories 2–7) 11
Category 2. Third-Party Transaction Monetization 12
Category 3. Own-Offer Monetization 13
Category 4. End-User Surplus Extraction 15
Category 5. Sponsored Recommendation 16
Category 6. Advertising Inventory Generation 19
Category 7. Preference Learning 19
C. Objectives Aimed at Future Value Creation (Categories 8–12) 20
Category 8. Engagement and Retention 20
Category 9. Perception Shaping 21
Category 10. Operational Demand Steering 22
Category 11. Ecosystem Dependence 23
Category 12. Preserving the Gateway 25
D. Influence, Welfare, and Integrity (Categories 13–16) 27
Category 13. Third-Party Persuasion 27
Category 14. Ideological and Social Influence 28
Category 15. Welfare and Public Interest 28
Category 16. Platform Integrity and Safety 30
E. Design Levers and Recommendation Objectives: Cross-Cutting Observations 31
1. The Inferential Gap Between Design Levers and Objectives: Counterfactuals, Incentives, and Effects 32
2. Bridging the Inferential Gap Between Design Levers and Objectives 33
3. Regulatory Intervention Under Residual Uncertainty About Platform Objectives 35
IV. RECOMMENDATION POWER UNDER THE DMA 35
A. On the DMA's Recommendation Vocabulary 36
B. Recommendation Power in the DMA's Designation Framework 37
1. The Intermediation-Power Component 37
2. The Information-Power Component 38
C. DMA Obligations Constraining Recommendation Design 40
1. Article 6(5) DMA: The Self-Preferencing Prohibition 40

2. Beyond Self-Preferencing: Other DMA Constraints on Recommendation Design ... 42
3. Mapping the Taxonomy onto the DMA ... 44
4. Article 13 DMA: Recommender Design as Circumvention ... 46
V. CONCLUDING REMARKS ... 49
REFERENCES ... 51

I. INTRODUCTION

Attention is scarce while options are not. On digital platforms, the number of products, services, sellers, apps, videos, or accommodations that could in principle be shown to a user vastly exceeds what any user can inspect.[1] Platforms therefore allocate scarce attention: through ranking and filtering, through recommendations and feeds, through defaults and preinstallation, and through the prominence given to particular offers, labels, and modules. They may also shape how presented options are perceived. Whoever controls the presentation of options also influences which options users notice, compare, and choose—and hence which transactions occur, which businesses grow, and which content is consumed.

Two established concepts capture the underlying forms of control. *Information power* denotes the control that information intermediaries acquire by selecting, ordering, and ranking the options that users see, thereby governing what users notice, compare, and choose. *Intermediation power* denotes the power that follows from a platform's position between market sides, in particular where business users depend on the platform to reach end users. This article draws out *recommendation power* as the point where the two meet: the platform steers user choice while allocating the scarce visibility on which business users' effective access depends.

Why does the concept matter? Two related observations motivate the analysis that follows. First, recommendation power strains ex post abuse control. Ranking and recommendation are continuous design choices rather than discrete exclusionary acts; users and business users often cannot observe whether prominence reflects relevance, payment, self-preferencing, or strategic steering; and the harm may be long-lasting—lost visibility can mean lost scale, data, reputation, and future contestability, so that, by the time ex post remedies arrive, the competitive process may already have shifted.[2] Second, those same features help explain why recommendation power is relevant to the ex ante regime of the Digital Markets Act (DMA),[3] which seeks to ensure contestable and fair markets in the digital sector where gatekeepers are present.[4] Control over visibility bears on contestability[5] when it entrenches the core platform service (CPS) as a gateway or impedes undertakings seeking to enter, expand, or challenge the gatekeeper. It also bears on fairness[6] when it enables the gatekeeper to discriminate against dependent business users or extract disproportionate advantages from them.

The article makes three contributions in its three substantive sections. First, it articulates the concept of recommendation power and locates it in the intellectual trajectory that leads from the analysis of information intermediaries and of intermediation power to the DMA and equivalent regulatory instruments, such as Section 19a of the German Competition Act[7] (Section II).

[1] On digital attention intermediaries and the mechanisms by which they attract and allocate consumer attention, see Peitz (2024).
[2] See below Section II.D. See also Crémer, de Montjoye, & Schweitzer (2019), 39–54, 66–70; Schweitzer, Haucap, Kerber, & Welker (2018), 66–78, 93–105; Schweitzer & Welker (2019).
[3] Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 Sept. 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act), O.J. 2022, L 265/1 (DMA).
[4] Article 1(1) DMA.
[5] Recital 32 DMA.
[6] Recital 33 DMA.
[7] Gesetz gegen Wettbewerbsbeschränkungen (Act Against Restraints of Competition).

Second, the article proposes a taxonomy of the objectives that guide platform operators' recommendation design (Section III). By "objectives" we mean the objective functions embodied in the design, that is, the criteria that explain how and why users' attention is directed in particular ways. The taxonomy distinguishes an independent-advice benchmark from fifteen other sixteen objectives that may motivate systematic departures from it. It classifies *objectives*, not *design levers*: the same lever—self-preferencing, sponsored slots, personalization, filtering, defaults, suppression—can serve different objectives. The taxonomy is descriptive. Classifying an objective carries no normative implications for how the law should be designed, interpreted, or applied. It nevertheless lays the analytical foundations for subsequent normative analysis by providing a framework within which different regulatory responses can be assessed. Throughout, we use a number of competition cases in which ranking, recommendation, default placement, dispatch, rating signals, or visibility allocation formed part of the alleged anticompetitive conduct as illustrations of particular objectives and mechanisms within the taxonomy.[8]

Third, the article analyzes how the DMA constrains recommender design (Section IV). The DMA does not regulate recommendation power as such, nor does it impose a general prohibition on biased recommendations by designated gatekeepers. Instead, it targets specific forms of bias and imposes other obligations that constrain recommendation design. These obligations include direct rules governing ranking (most prominently, the prohibition of self-preferencing in ranking under Article 6(5) DMA), as well as rules governing data use and transparency, all of which shape recommendation design. Moreover, Article 13 DMA serves as the Regulation's anticircumvention and future-proofing provision. It also extends as far as recommendation design that reproduces the effects of practices prohibited elsewhere in the DMA, whether or not those practices themselves concern steering. De-ranking as a substitute for prohibited price-parity clauses is the paradigmatic example.[9]

The article proceeds as follows. Section II develops the concept of recommendation power. Section III presents the taxonomy of recommendation objectives. Section IV analyzes recommendation power under the DMA: the DMA's vocabulary of ranking and recommendation, Article 6(5) DMA and its economic rationale (including an unintended potential downside of its enforcement), the provisions beyond Article 6(5) DMA, a mapping of the taxonomy onto DMA obligations, and the role of Article 13 DMA. Section V concludes.

[8] These cases include the EU *Google Shopping* case (note 43), the *Amazon* proceedings concerning the Buy Box, FBA and/or related marketplace practices in the EU (note 81), Italy (note 80), Germany (note 87), and the UK (note 82), the FTC's pending case against Amazon (notes 59 and 88), the South Korean cases against Coupang (note 44), Kakao Mobility (note 74), and Google (note 89) as well as the Japanese *Google* case (note 89) and the German *netdoctor/Google* case (note 96). For the underlying decisions and their procedural status, see the references accompanying the respective illustrations in Section III.

[9] Franck & Peitz (2024). The present article builds on the effects-based, three-step anticircumvention methodology developed there and applies it to recommendation design.

II. THE CONCEPT OF RECOMMENDATION POWER

A. Information Power

Information power arises where information intermediaries determine the set and ordering of the options that users notice, compare, and choose.[10] Search engines and comparison and rating portals collect and index content and, applying relevance criteria, filter and rank it; they are typically the user's first point of orientation. The exercise of this curatorial control is, in itself, the service: An unfiltered, unordered list would be useless to the user. Selection becomes power when users rely on the intermediary: A provider that is not shown is, for those users, not found. Ordering and ranking therefore govern not only discovery but also the allocation of users' attention and, through it, their transaction decisions.

B. Intermediation Power

Intermediation power is a platform's ability, by virtue of its intermediary position, to control effective access between business users and end users.[11] It increases with the degree to which business users depend on the platform to reach end users and it is strongest where end users single-home and cannot easily do without the intermediation service, so that business users have to go through the platform to reach customers at scale. Moreover, the relevant variable is not only access in a formal sense but the quality of access: listing, ranking, prominence, trust signals, and placement. A business user that is formally present on a platform but relegated to positions that users never reach does not enjoy effective access. In short, intermediation power concerns effective rather than formal access to end users.

C. Recommendation Power: the Intersection

Recommendation power arises where the exercise of information power is also an exercise of intermediation power. Information power allows an intermediary to use selection, ordering, and ranking to shape what users notice, compare, and choose and therefore which providers they find. Where business users also depend on that intermediary to reach sufficiently many end users, the allocation of visibility operates through an effective-access bottleneck. Rankings, recommendations, defaults, labels, prompts, feeds, and prominence then steer user choice while governing the quality of dependent business users' access to those users. Recommendation power thus denotes this particular conjunction of informational control and intermediation power: The visibility that a platform allocates becomes the scarce resource through which it controls the quality of business users' effective access to end users.[12]

---

[10] Peitz & Schweitzer (2016), 825, 828–829, emphasizing the role of search engines as information intermediaries that select among available information and thereby determine its visibility to users.

[11] For the underlying conception, see Schweitzer, Haucap, Kerber, & Welker (2018), 66–73; Schweitzer & Welker (2019); Crémer, de Montjoye, & Schweitzer (2019), 4 (linking "intermediation power" to the concept of "unavoidable trading partner"). The concept of intermediation power was central to the rationale underlying the introduction of digital gatekeeper regulation under the DMA, see European Commission, *Impact Assessment Report accompanying the Proposal for a Digital Markets Act*, SWD(2020) 363 final, Part 1, paras. 14, 27, 31, 143–144; see also pp. 56–57 ("ability to impose market rules due to its intermediary function"; "ability to misuse its intermediation position to its own competitive advantage").

[12] The concept of recommendation power developed here combines two strands identified in the literature: platforms' intermediation power, arising from business users' dependence on access to the other market side, and platforms' informational and rule-setting power exercised through ranking and other forms of selection. See Crémer, de Montjoye, & Schweitzer (2019), 54, 64–70; Schweitzer, Haucap, Kerber, & Welker (2018), 66–78; Schweitzer & Welker (2019).

Two clarifications help sharpen the concept. First, recommendation power is not confined to ranked lists. Defaults, preinstallations, buy boxes, badges, snippets, notifications, voice answers, and single-result interfaces serve the same function: They make one option salient at the expense of others.[13] As will be seen, the DMA's definition of "ranking" deliberately reflects this functional breadth.[14]

Second, recommendation power, as an analytical concept, is normatively neutral. The same is true of power more generally if understood as the capacity to bring about intended outcomes.[15] In this sense, recommendation power expands the feasible set through which a platform can pursue whatever objective function it pursues. This expanded set of strategic possibilities includes socially beneficial uses. Recommender systems are an information-push strategy: They reduce search costs, they can improve expected match quality, and they generate platform-specific network effects, since more users and transactions improve recommendation quality and make the platform more attractive; in two-sided settings, better recommendations create cross-group external benefits.[16] Yet, even if recommendation power can be exercised for beneficial purposes, a platform's ability to pursue its objectives by steering users' choices remains a phenomenon that merits close scrutiny. In sum, whether and under what conditions the exercise of recommendation power gives rise to policy concerns depends on the objectives it serves, which is the subject of the taxonomy in Section III of this article.

D. Platforms as Private Rule-Setters and the Limits of Ex Post Control

Recommendation power is not merely a mechanism through which platforms influence individual user choices. It also has an institutional dimension. Because recommendation systems determine which options users see, compare, and ultimately choose, they become a central instrument through which dominant platforms shape the rules of interaction between users and business users. This is precisely why Crémer et al.'s influential 2019 expert report for the European Commission characterized dominant platforms as private rule-setters.[17] Recommendation systems are a core rule-setting instrument: They determine the order in which options appear and decide which options are made salient. Seen through this lens, recommendation power is not merely influence over individual choices but authority over the terms of competition among business users: A platform that connects users and business users is also a private market designer that sets the rules on the platform. If platforms lack incentives to set fair, unbiased, and pro-user rules, regulatory intervention may be needed.

Three mutually reinforcing characteristics of recommendation power explain why traditional ex post abuse control often struggles. First, *continuous recalibration of design*: Ranking and recommendation are ongoing governance choices, not only discrete exclusionary acts; an abuse inquiry that isolates a single act at a single point in time sits awkwardly with a design parameter that is continuously adjusted.

[13] See ECJ, 2.7.2026, Case C-738/22 P, Google and Alphabet v Commission (Google Android), ECLI:EU:C:2026:533, paras. 203, 205–213, 217–219 confirming that preinstallation may generate a status quo bias notwithstanding users' ability to download competing apps, thereby systematically shifting actual user choice towards the preinstalled option.
[14] Article 2(22) DMA and Recital 52 DMA; see below Section IV.B.
[15] As reflected in Russell's classic definition (1938, 35) ("Power may be defined as the production of intended effects").
[16] Belleflamme & Peitz (2021), ch. 6. The same instrument that creates value can become a source of market power, and market power can be abused.
[17] Crémer, de Montjoye, & Schweitzer (2019), in particular 5, 61–62, 71. See also Eifert, Metzger, Schweitzer, & Wagner (2021), 991, 994–995.

Second, *opacity*: Users and business users may not observe whether prominence reflects relevance, payment, self-preferencing, or strategic steering; the informational asymmetry runs in favor of the platform, which controls both the algorithm and the data needed to evaluate it. Third, *dynamic harm*: Lost visibility can mean lost scale, data, reputation, and future contestability. The harm therefore accrues long before a literal refusal to deal, and it accrues while proceedings under Article 102 TFEU are still pending; by the time ex post remedies arrive, the competitive process may already have shifted.[18]

These features help to explain the road to the DMA. The DMA targets undertakings whose CPSs are important gateways for business users to reach end users.[19] The Commission's impact assessment ahead of the DMA placed visibility and intermediation side by side: It identified an "important intermediation function" and an "important 'visibility' role" as characteristic features of the CPSs in which problematic gatekeeper practices are particularly severe. It also directly recognized biased recommendations—platforms may favor their own services or bias recommendations toward a content provider that charges lower royalties—and, most prominently, treated self-preferencing in ranking as capable of undermining contestability and fairness, not merely as a ranking-design choice.[20] The logic can be stated in one line: Where visibility control coincides with user dependence, ex ante obligations to protect contestability and fairness become justified.

The German legislature had already set in motion a related regulatory response before the Commission proposed the DMA. Section 19a of the Competition Act, which entered into force in January 2021, addresses undertakings of paramount significance for competition across markets and is designed to deal with ecosystem and gatekeeper power beyond classical dominance concepts.[21] The legislative materials emphasize a new form of market power rooted in a key intermediary position between multiple market participants, data advantages, and the ability to shape the rules of interaction and personalize offers. Correspondingly, the prohibitions target self-preferencing and access foreclosure in ways that map directly onto recommendation design, including presentation, ranking, listing,

[18] Crémer, de Montjoye, & Schweitzer (2019), 31, 35–36; Schweitzer, Haucap, Kerber, & Welker (2018), 12, 59–62; Schweitzer & Welker (2019).
[19] Article 3(1)(b) DMA. For details, see Section IV.B.1.
[20] European Commission, *Impact Assessment Report accompanying the Proposal for a Digital Markets Act*, SWD(2020) 363 final, Part 1, paras. 41–43, 50 (expressly identifying an "important intermediation function" and an "important 'visibility' role" as characteristic features of CPSs).
[21] The government draft of the Tenth Amendment to the Competition Act of October 2020, introducing Section 19a (Deutscher Bundestag, Drucksache 19/23492, 19 October 2020, Gesetzentwurf der Bundesregierung, Entwurf eines Gesetzes zur Änderung des Gesetzes gegen Wettbewerbsbeschränkungen für ein fokussiertes, proaktives und digitales Wettbewerbsrecht 4.0 und anderer wettbewerbsrechtlicher Bestimmungen), predates the Commission's DMA proposal of December 2020, and Section 19a entered into force more than a year and a half before the DMA was adopted by the European Parliament and the Council in September 2022. By moving ahead, the German legislature put pressure on the EU to address the digital gatekeeper challenge and arguably created an additional impetus for the DMA: The prospect of diverging national gatekeeper regimes strengthened the case for a harmonized internal-market framework as reflected in the DMA's legal basis in Article 114 TFEU. See Franck & Peitz (2021), 226–227, on the risk of regulatory fragmentation in the internal market. This is not to say that Section 19a of the Competition Act has been superseded by the DMA: as national competition law, its application to gatekeepers is preserved under Article 1(6)(b) DMA insofar as it amounts "to the imposition of further obligations." For the legal framework defining the complementary role of Section 19a of the Competition Act, as well as the scenarios in which that role may become practically relevant, see Franck (2025), 230–233.

preinstallation, integration, and alternative access routes.[22] While Section 19a of the Competition Act reflects the same underlying structural concern, the remainder of this article focuses on the DMA.

E. Summary

To summarize the conceptual part: Information power explains how platforms shape users' decision environment; intermediation power explains why business users depend on platforms for effective access to end users; recommendation power captures the exercise of the former through the latter, as rankings, recommendations, defaults, curation, prominence, labels, and interface design allocate the scarce visibility on which effective access depends. The latter concept thus provides a common analytical framework for understanding Crémer et al.'s report, the DMA, and Section 19a of the German Competition Act as responses to the same structural problem: Where access to users runs through the platform, what the platform makes prominent becomes what is chosen, and control over prominence becomes control over access.

III. A TAXONOMY OF RECOMMENDATION OBJECTIVES

The concept of recommendation power explains *how* platforms steer users' choices. It does not yet explain *why* they do so. This explanation lies in the objective functions embodied in recommendation design. This section develops a taxonomy of recommendation objectives that classifies these objective functions. It provides the analytical bridge between the descriptive concept of recommendation power developed in Section II and the legal analysis in Section IV of how the DMA constrains recommendation design.

A. What Is Being Classified—and Relative to Which Benchmark?

The taxonomy classifies the objectives that platforms may pursue through recommendation design. Each of the following categories identifies a distinct objective that a platform may pursue through recommendation design. Recommendation bias arises where recommendation design systematically departs from the independent-advice benchmark (category 1) in order to pursue one of the objectives numbered here as 2 to 16. The same design lever—self-preferencing, sponsored slots, personalization, ranking, filtering, defaults, suppression—can serve several objectives, and the same objective can be pursued through different levers.

The taxonomy is descriptive in the sense that observing a particular objective in recommendation design does not, on its own, imply that the objective is socially beneficial or harmful, or that it does or does not warrant regulatory intervention. Each category asks what the platform seeks to maximize or protect (revenue, attention, data, retention, dependence, welfare, safety). The objective is pursued through design levers (ranking, labels, defaults, filters, sponsored slots, notifications) and manifests itself in observable outcomes (exposure, traffic, conversion, engagement metrics, trust measures). Each category therefore illustrates the objective by reference to recommendation designs and practices through which it may be pursued. Because the same design may serve several objectives—including objectives beyond those captured by the taxonomy—the examples are illustrative rather than conclusive

---

[22] For a detailed overview of the conduct covered by Section 19a of the Competition Act, see Franck & Peitz (2021), 519–525.

and do not establish that a platform in fact pursues the objective in question. Section III.E returns to the question of how such objectives may be inferred from observable recommendation practices.

Before turning to the various categories, one further point requires clarification. Recommendation bias is inherently benchmark-dependent because any analysis of recommendation bias requires an explicit counterfactual: Inferior relative to what? Category 1 shows that the benchmark of independent advice is itself not a single standard but a family of neutral benchmarks, which we spell out below.

Category 1. The Independent-Advice Benchmark

The taxonomy begins with the benchmark against which the remaining categories are assessed. Under the independent-advice benchmark, recommendation design serves the users alone. The platform still selects and ranks—selection and ranking are the very service it provides—but no objective other than improving the user's decision shapes the recommendation. Examples include neutral product comparison, relevance-based search, tariff or plan comparison based on user benefits, and matching designed solely to maximize expected end-user welfare.[23] In the ideal case, the intermediary finds the best match from the user's perspective, allowing the end user to rely on the recommendation without independently evaluating the available alternatives.[24]

Speaking of "the" benchmark oversimplifies matters. The intermediaries considered here simultaneously serve multiple user groups, most importantly end users and business users. Independent advice therefore does not refer to a single benchmark but to a family of neutral benchmarks. Three candidates are natural. Under an *end-user-optimal* benchmark, the ranking maximizes expected end-user surplus—the net match value to users, which is the difference between gross match value[25] and price. This is a particular recommendation design. More generally, recommendations may be made according to a score that places more weight on either gross match value or price dimension.[26] Under a *business-user-optimal* benchmark, the ranking gives suppliers neutral, nondiscriminatory exposure according to relevance or quality—or, more demandingly, maximizes business-user surplus. Under a *joined user* benchmark, the ranking induces the transactions and matches with the highest joint surplus of end users and business users; more generally, any weighted sum of the two defines a benchmark. Reimers and Waldfogel (2023) turn this into an operational definition: A platform's ranks are biased if the outcomes they generate lie below the frontier that maximizes a weighted sum of consumer and seller surplus. Bias is then a distance from a welfare frontier rather than a departure from a unique "correct" order. The distinction matters because the three benchmarks may point to different rankings. The ranking that maximizes end-user surplus may concentrate exposure on a few best matches and thus deny "fair" exposure to most business users, while exposure allocated evenly or proportionally across business users need not serve users well. The same observed ranking can therefore be biased under one

[23] Bakos (1997); Dinerstein, Einav, Levin, & Sundaresan (2018); Belleflamme & Peitz (2021).
[24] Peitz (2024).
[25] Gross match value is the value an offer provides to a particular user before taking its price into account.
[26] Formally, recommendations may be based on a weighted difference between gross match value and price. If the respective weights are denoted by λ and $(1-\lambda)$, the recommendation score is $\lambda v-(1-\lambda)p$. At $\lambda=1$, recommendations are made by gross match value; at $\lambda=1/2$, by net match value; and at $\lambda=0$ exclusively by price. Rhodes & Zhou (forthcoming) study a search market with competing sellers and heterogeneous privacy costs: consumers' disclosure of personal data allows the platform to make personalized recommendations and the weight given to price in the recommendation design affects consumers' disclosure decisions, price levels, and consumer welfare.

benchmark and unbiased under another. What all members of the family share—and what makes them *neutral*—is that the platform's own surplus receives no weight; categories 2 to 16 describe what happens when it does.

Each benchmark, moreover, comes in variants depending on what is treated as fixed and what is treated as endogenous, that is, as responding to the ranking policy chosen. The first is the *supply-taken-as-given* variant (ranking by user fit). Here, the set of participating business users and their prices, qualities, and content are fixed, and the ranking is a pure matching device: It allocates given offers to users according to fit. This is also the natural setting for empirical bias detection, which conditions on the observed choice set and observed prices.[27] In the *equilibrium* variant (ranking by user value), business-side conduct becomes endogenous. Sellers adjust prices and quality, and creators adjust content choices in anticipation of the ranking policy; sellers and creators may also enter or exit. This variant internalizes those responses. An end-user-optimal policy of this kind systematically rewards low prices and high quality and thereby disciplines the business side. Nor is the end-user side entirely exogenous: How much users rely on recommendations (and hence how attention is allocated) depends on the anticipated informativeness of the advice, so that better recommendations attract more attention and generate recommender-driven network effects.[28] The distinction between the variants is not merely academic: As Section IV.C will show, the equilibrium, end-user-value variant may be the design most exposed to legal challenge under Article 6(5) DMA.

B. Objectives Aimed at Immediate Value Capture (Categories 2–7)

The first group comprises objectives that use recommendation design for immediate value capture. Value is captured through transactions by earning business-user commissions (category 2), increasing the platform's own product margins (category 3), better extracting end users' willingness to pay (category 4), the sale of prominence (category 5), or the creation of monetizable resources, namely attention (category 6), and data (category 7). This distinguishes the first group from the second group of objectives, discussed in Section III.C, where recommendation power is used not for immediate value capture but to increase future earning opportunities.

The objectives grouped together in this first group are not mutually exclusive and may reinforce one another. Accordingly, recommendation design that serves one objective will often advance another at the same time. For example, a design that maximizes attention also generates data, while recommendation design that increases transactions may simultaneously increase attention.

Within the first group, the first three objectives form a triptych because they all seek to capture value through transactions. They differ only in the variable on which recommendation design is conditioned: business-user margins (category 2), platform's margin as a seller (category 3), or end-user willingness to pay (category 4).

The last two objectives, attention and data, play a dual role within our taxonomy. They may be monetized immediately, for example through advertising, but they also generate future value. Where the delayed payoff, rather than current monetization, becomes the primary objective, the conduct falls within the second group—"future value creation"—discussed in Section III.C.

[27] Reimers & Waldfogel (2023).
[28] Peitz (2024); Belleflamme & Peitz (2021).

Category 2. Third-Party Transaction Monetization

This objective seeks to capture value by increasing the platform's net revenue from recommended third-party transactions, matches, leads, or content consumption. Depending on the business model, recommendation design may optimize conversion, commission revenue, lead payments, basket size, or net contribution after royalties. Recommendation departs from the independent-advice benchmark when the more profitable option is not the best match. Users see a ranked list, a buy box, a default, or a "recommended" label; the platform sees the profit associated with each option. Three mechanisms explain how recommendation design can become conditioned on that profit rather than user benefit.

The first mechanism arises from *fee heterogeneity*: The platform earns different net returns across third-party offers—differential commissions, lead payments, or royalty costs. Where recommendation design allocates prominence according to platform profit rather than user benefit, offers carrying higher commissions, lead payments, or lower royalty costs tend to systematically receive preferential treatment. The direction of harm is not mechanical, however. Whether biased intermediation harms users depends on the extent to which platform revenue per transaction aligns with the surplus users derive from that transaction. Where platform revenue and user surplus are aligned, bias is small; where they conflict, steering diverts users to worse matches.[29] The misalignment is starkest where the platform monetizes visits or leads rather than completed transactions: It may then profit from diverting search by deliberately presenting less-suitable matches first to generate additional monetizable visits.[30] A closely related variant of the mechanism arises on content platforms: A streaming service that pays per-stream royalties earns more, per unit of listening or viewing, from content with lower royalty rates and therefore has an incentive to tilt recommendations toward such content—monetization by cost avoidance rather than by revenue.[31] The mechanism also affects competition among business users. Prominence sold against the highest commission feeds back into seller conduct: Sellers compete for the recommendation slot through fees rather than through price or quality, and equilibrium prices can rise as commissions are passed through.[32]

Unlike the first mechanism, the second operates even under *uniform ad valorem fees*. If the platform takes the same revenue share from every transaction, its revenue per recommendation is proportional to the transaction price multiplied by the probability of purchase. If purchase probabilities are held constant, recommendation then tilts toward higher-priced options and larger baskets, whereas user benefit depends on surplus, not expenditure. The role of the uniform fee is important. Because the fee itself cannot discriminate between sellers, differences in the platform's take arise solely from differences in price. Recommendation design therefore becomes the instrument through which demand is steered toward higher-priced offers. Whether this conflicts with end-user benefit depends on what price differences reflect: Where a higher price corresponds to a higher end-user surplus, revenue-based and user-based rankings are aligned. Where they are not, the platform gains from demoting those offers that would generate the greatest benefit for end users. A closely related mechanism operates through conversion-rate ranking. When a low conversion rate signals poor fit or low quality, ranking by conversion serves users. But conversion also falls when users discover a product on the platform and

[29] de Cornière & Taylor (2019), who provide a general framework in which the welfare effect of biased intermediation depends on whether the intermediary's and the users' payoffs are congruent or conflicting.
[30] Hagiu & Jullien (2011).
[31] Bourreau & Gaudin (2022).
[32] Armstrong & Zhou (2011).

purchase it more cheaply elsewhere; a conversion-based ranking then systematically demotes offers whose off-platform counterparts are attractive to end users. What began as a relevance proxy becomes a mechanism for penalizing leakage.[33] We will return to this pattern under gateway preservation (category 12) and with regard to the DMA's anticircumvention analysis.[34]

The third mechanism arises where end users are fallible and therefore prone to decision errors. Offers that exploit consumer errors, for example through hidden charges, drip pricing, or contract terms whose costs are underestimated, may convert particularly well and therefore yield high platform revenues. Recommendation design that maximizes conversion or revenue consequently treats consumer errors as a signal of match quality, amplifying rather than correcting behavioral distortions.[35]

These mechanisms share a common structure: Each creates a wedge between the platform's payoff from a recommendation and the benefit that recommendation generates for the end user. The concern is greatest where users perceive recommendations as relevance signals while business users depend on the resulting traffic. Examples include steering users toward high-commission sellers, affiliate links, lead generation, cross-selling, lower-royalty streaming content, and the demotion of offers prone to showrooming or bypass.

Category 3. Own-Offer Monetization

The objective seeks to capture value by increasing revenue from the platform's own offers, whether through direct sales, subscriptions, or service fees. Direct self-preferencing belongs to this category where the platform promotes its own offer. Unlike category 2, where the platform monetizes third-party transactions as an intermediary, the platform here monetizes its own products or services as the supplier. The underlying wedge is a margin comparison: On the own offer, the platform earns the full retail or subscription margin; on the recommended third-party offer, it earns only the intermediation fee. Whenever the own margin exceeds the fee that a third-party sale would generate, the platform may have an incentive to allocate prominence to its own offer, even where, at given prices, users would prefer the third-party option. A cost-side variant arises on content platforms: A vertically integrated service avoids paying royalties on its own content, so recommending a platform's own content is monetization by cost avoidance—the analogue of the low-royalty steering discussed under category 2.[36] Examples include the promotion of a platform's own private-label products, proprietary apps, logistics services, payment products, and media content.

Why does the platform bias prominence rather than simply raise its fees until third-party sales are equally profitable? The answer lies in the fee instrument being subject to a constraint. First, some recommendations carry no fee by design, organic search results being the leading example. Because the platform cannot appropriate any of the value created by a third-party match, the incentive to divert traffic to its own monetizable offer is strongest. Second, even where fees can be charged, they are often uniform across broad product or service categories. Such fee structures may be adopted for reasons of administrability, to limit discrimination disputes, or to preserve sellers' investment incentives by reducing hold-up concerns. Because the fee cannot be tailored to individual products, the platform

[33] Hunold, Kesler, & Laitenberger (2020); Hagiu & Wright (2024) on platform responses to marketplace leakage.
[34] See below Section IV.C.4.
[35] Heidhues, Köster, & Kőszegi (2023).
[36] On integrated platforms' incentives to bias recommendations toward own content whose distribution avoids third-party payments, see Peitz (2024); on the general margin logic, de Cornière & Taylor (2014).

cannot participate in the higher margins of some third-party offers. Recommendation design therefore again becomes the substitute instrument.[37] This incentive logic provides an economic rationale of Article 6(5) DMA, to which we return in Section IV.C.

Prominence for the platform's own offer appears benign where that offer genuinely provides the best match, and vertical integration may even strengthen incentives to improve quality. By contrast, it appears harmful where an inferior own offer displaces the preferred third-party alternative.[38] The welfare assessment is, however, more complex. Theory identifies conditions under which vertically integrated platforms have an incentive to favor their own offers, and empirical work documents that such own-offer advantages occur in practice.[39] The existence of that incentive, however, does not by itself establish a welfare harm. Own-offer prominence generally results in partial rather than complete foreclosure: Third-party offers remain available but lose traffic, scale, and data. Whether this ultimately harms welfare depends on equilibrium responses on both sides of the platform. The platform's own-product entry may discipline third-party prices and increase variety, while seller participation and entry adjust to the anticipated ranking policy and fee structure.[40]

As a result, an equilibrium-optimal ranking may be preferable from an end-user perspective to a simpler ranking that is optimal only when taking supply as given, precisely because the former anticipates how sellers respond to the prospect of greater visibility and uses ranking to induce lower prices, higher quality, or better realized matches. Such a design uses visibility actively as an incentive instrument: offers receive greater prominence not merely because they are currently more attractive to consumers but because rewarding particular seller choices may induce responses that benefit end users. Accordingly, what may therefore appear at first sight to be a preferential recommendation design aimed at own-offer monetization (category 3) may instead form part of a strategy that ultimately serves end users' interests (category 1).

Admittedly, while antitrust enforcement should in principle remain open to such considerations, establishing them in an individual case is likely to be difficult. In particular, it may be hard to determine whether observed self-preferencing is genuinely an intermediate step toward an end-user-optimal equilibrium rather than a means of monetizing the platform's own offering. Empirically establishing the relevant counterfactual and equilibrium effects may require considerable effort.[41] Under the

---

[37] Peitz (2022); Franck & Peitz (2024), section 5.2.1; Feasey & Krämer (2019). Uniform fees may be preferred for administrability, to reduce exposure to discrimination complaints, and to mitigate hold-up vis-à-vis third parties making platform-specific investments.

[38] The dynamic consequences of such diverted traffic (rivals' loss of scale, data, and future contestability) are potentially important, but within the taxonomy they mark the transition to different objectives: where diminishing rivals is itself the aim, the conduct falls under preserving the gateway (category 12) and, where the traffic is steered for the data it generates, under preference learning (category 7).

[39] de Cornière & Taylor (2014) for the theory of integration and search engine bias. Chen & Tsai (2024) find that Amazon's "Frequently Bought Together" recommendations favor Amazon-sold products—exploiting Amazon stockouts for identification—in a pattern explained by steering incentives rather than consumer preference; Farronato, Fradkin, & MacKay (2023) document the prominence of Amazon's own brands in search rankings.

[40] Hagiu, Teh, & Wright (2022) analyze the equilibrium of dual-mode platforms in which fees, first-party presence, and seller participation are jointly determined; Etro (2023) endogenizes seller entry onto the platform; Zennyo (2022) studies platform encroachment and own-content bias. This literature also evaluates structural counterfactuals and finds their welfare effects ambiguous, as constrained platforms adjust fees—a response that recurs, in the context of the DMA's obligations, in Sections IV.C and IV.F.

[41] Peitz (2022); Lee & Musolff (2026); Farronato, Fradkin, & MacKay (2023); Reimers & Waldfogel (2023).

prohibition of self-preferencing in ranking in Article 6(5) DMA, by contrast, such an "equilibrium" justification would appear to be excluded by the nature of the regulatory approach itself.[42]

Two prominent enforcement actions illustrate the category of own-offer monetization. In *Google Shopping*, the Commission found that Google gave its own comparison shopping service prominent placement in general search, while rival services appeared only as generic search results ("blue links") and were subject to demotion algorithms. According to the Commission's theory of harm, this discriminatory self-preferencing leveraged Google's dominance in general search into comparison shopping through diverting traffic and reducing the visibility of rivals. The infringement decision and the fine were subsequently upheld on appeal.[43] *Coupang* illustrates the same objective in a different setting. The Korea Fair Trade Commission (KFTC) found that Coupang, which operates a vertically integrated online marketplace, reserved fixed high-ranking positions for private-brand and direct-purchase items and used employee reviews to inflate review counts and average ratings. According to the KFTC, the resulting artificial prominence and inflated ratings misled consumers, diverted demand away from third-party sellers, and weakened incentives for price competition. The case also illustrates that own-offer monetization can operate not only through ranking positions but equally through trust signals such as ratings and reviews, thereby connecting this category with perception shaping (category 9).[44]

Category 4. End-User Surplus Extraction

This objective seeks to extract surplus from end users by steering them across products, qualities, bundles, or purchase decisions according to their willingness to pay. Rather than discriminating through prices, the platform discriminates through recommendation design, shaping users' perceived match values while addressing heterogeneity across end users. The defining feature is that, where prices are uniform or otherwise constrained, recommendation design may become the instrument of discrimination. Surplus is then extracted through what end users are shown rather than through what they are charged.[45] Unlike categories 2 and 3, this mechanism requires no divergence between the platform's interests and those of business users. It operates equally where the platform is vertically integrated and where its incentives are perfectly aligned with those of the business-user side. The wedge runs solely between the platform and end users.

The mechanism of inflated recommendations illustrates the logic. Consider an experience good. Consumers cannot assess their match value before purchase, whereas the seller-recommender observes

---

[42] See below Section IV.C.1., text accompanying fn. 143.

[43] 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, 11.10.2024, Case C-48/22 P, *Google and Alphabet v Commission (Google Shopping)*, ECLI:EU:C:2024:726. For the standard ultimately developed by the Court of Justice in this case for when self-preferencing constitutes an abuse under Article 102 TFEU, see, e.g., Fischer, Hornkohl, & Imgarten (2025).

[44] Korea Fair Trade Commission (KFTC), 5.8.2024, Decision No. 2024-284, Case Nos 2021Seogam1429 and 2022Seogam0713, *Coupang and CPLB*; see, for an English-language account of the decision, KFTC, 13.6.2024, Press Release, "KFTC Imposes Severe Sanctions on Coupang for Consumer Deception." According to media reports, Coupang challenged the decision before the Seoul High Court; proceedings on the merits were pending at the time of writing. See Kang Min-woo & Kim Si-kyun, "Court Suspends KFTC Corrective Order against Coupang over Alleged Search-Ranking Manipulation," *Maeil Business Newspaper*, 10 October 2024 (in Korean).

[45] Peitz & Sobolev (2025, 2026a, 2026b). The information-design perspective goes back to Kamenica & Gentzkow (2011): The recommender commits to a policy that shapes consumers' posterior beliefs.

it and commits to a recommendation policy. Suppose that prices are uniform and exceed the unconditional expected valuation. In that setting, informative recommendations become essential because they help consumers to learn whether the product is in fact worth buying. A profit-maximizing recommendation policy, however, is not fully informative. Instead, it sometimes recommends purchase even to consumers whose valuation falls short of the price, pooling them with consumers for whom the purchase is genuinely worthwhile. Such recommendations reduce expected consumer surplus and are therefore biased under the end-user benchmark in category 1. Where consumers' valuations also fall below production cost, the recommendations additionally induce inefficient transactions and become biased under the total-surplus benchmark as well.[46]

Price and recommendation policy are closely linked. The more recommendations are inflated—that is, the more often they encourage purchase beyond what an unbiased recommendation would advise—the lower the price must be for consumers with a recommendation to remain willing to follow them. The seller therefore faces a trade-off between broad but inflated recommendations supported by lower prices and accurate, informative recommendations that sustain high prices.[47]

Two qualifications delimit the mechanism. First, the mechanism disappears where consumers are homogeneous and prices can adjust freely. In that case, the seller extracts surplus through the price rather than through recommendation design. Recommendation bias survives endogenous pricing only where consumers are heterogeneous in their tastes or in the information they possess ex ante, so that a single price must serve heterogenous consumers. Biased recommendations are then a partial substitute for price discrimination.[48] Second, the platform-optimal recommendation policy does not necessarily recommend "too much." Where some consumers know their match value while others rely on recommendations, the optimal policy depends on both the distribution of valuations and the seller's pricing incentives. Depending on these conditions, recommendations may be excessive or socially insufficient. Indeed, some consumers may receive no recommendation even though positive gains from trade exist. From a total-surplus perspective, recommendation bias therefore does not necessarily consist in "too much steering." Depending on the circumstances, it may equally consist in too little steering. Accordingly, an intervention that mandates fully informative recommendations cannot be assessed in isolation. Prices adjust in response, and whether end users and society ultimately benefit depends not only on the informativeness of recommendations but also on the induced price response.[49]

Category 5. Sponsored Recommendation

This objective seeks to capture value by directly monetizing prominence, that is, a privileged form of visibility, including ranking positions, recommendation labels, and other forms of endorsement-like placement. This sponsored recommendation differs from ordinary advertising in that the platform sells

[46] Peitz & Sobolev (2026a, 2026b). The two benchmarks can diverge: A policy may raise total surplus through better matching yet reduce consumer surplus through higher prices.
[47] The distortion does not rest on behavioral irrationality: Rational consumers follow the recommendation because, given the information structure the recommender controls, following it is optimal.
[48] The uniform-price constraint is essential: Rayo & Segal (2010) show that a monopolist who can condition the price on the signal sent to consumers optimally discloses fully; surplus is then extracted through pricing rather than through pooling. See also Roesler & Szentes (2017) on the set of outcomes attainable under monopoly pricing for arbitrary information structures, and Ennuschat (2025) for an analysis with partially informed consumers.
[49] Peitz & Sobolev (2025), who study a monopolist facing one segment informed about its valuation and another relying on the recommendation; see also the extension in Peitz & Sobolev (2026a).

placement *within the recommendation* itself. What the business user purchases is not merely exposure but the appearance of having been selected. The value lies in the prominence that end users may interpret as advice, quality, or a particularly good match. The underlying mechanism should be distinguished from those in categories 2 and 3. There, the platform biases a nominally organic ranking toward the offers from which it derives greater value. Here, by contrast, prominence itself becomes the product sold to business users, one recommendation slot at a time. Examples include sponsored search results, promoted marketplace listings, paid app placements, and native recommendation cards.

Selling prominence does not necessarily bias recommendations. Whether it does depends on how the prominence-allocation mechanism translates business users' willingness to pay into prominent placement and how this relates to user benefits. The first mechanism concerns how selling prominence *selects* who becomes prominent. A business user's willingness to pay for a prominent slot reflects the profit it expects from the resulting traffic (conversion multiplied by margin) and the pricing instrument determines which business users receive prominence, whether through an auction or through a posted per-click price. Where conversion tracks match quality, willingness to pay sorts business users by relevance: The market for prominence then aggregates dispersed information about match quality that the platform's algorithm may not possess, and end users may rationally inspect sponsored results first.[50] The same alignment conditions as in category 2 apply, however: Where business users' profit from prominence does not track end-user valuation, the market sorts on profitability rather than on end-user value. The design of the allocation mechanism can partially restore that alignment. Reserve prices, quality weights, and the level of the per-click price can screen out low-relevance bidders whose presence would waste end-user attention.[51]

The second mechanism concerns the interaction between sponsored and organic recommendations. Selling prominence can weaken the organic alternative in two ways. First, a sponsored position may displace an organic result and thereby reduce comparison.[52] Second, because organic recommendations are a free substitute for paid visibility, the platform may reduce organic quality.[53] Degradation can even take the form of deliberate dilution. By admitting low-relevance offers into the organic pool, the platform makes continued comparison unattractive for end users. Reduced comparison in turn relaxes price competition among the well-matched business users, whose increased profits the platform can subsequently tax through the price of prominence.[54] The presence of organic listings feeds back into bidding: Business users with strong organic positions bid differently from those without, so the composition of the sponsored slots and the platform's revenue depend on the organic ranking policy.[55]

---

[50] Chen & He (2011), where willingness to pay for placement increases with relevance, so that bidding sorts business users and paid placement becomes informative. Eliaz & Spiegler (2011) studies the same design problem when the platform sets a per-click price rather than running an auction, the price acting as a relevance screen on the pool of business users—the profit-maximizing design need not implement the relevance level end users would choose. See also Athey & Ellison (2011) on position auctions when consumers search rationally from the top.
[51] Athey & Ellison (2011) on reserve prices as a screen that economizes on consumer search costs.
[52] Motta & Penta (2026), where a business user may buy the leading sponsored position for its own offer in order to prevent a rival's organic result from appearing.
[53] White (2013) on the trade-off between organic ("left-side") quality and advertising ("right-side") profits; Burguet, Caminal, & Ellman (2015), showing that competition between a search engine's sponsored search advertising and publishers' display advertising may induce the search engine to reduce the reliability of organic content search; Belleflamme & Peitz (2021).
[54] Eliaz & Spiegler (2011), who show that a monopoly search engine may deliberately set a low per-click price so that low-relevance business users enter the pool, relaxing price competition among relevant ones.
[55] Xu, Chen, & Whinston (2012).

Recent work sharpens this logic: A platform that monetizes the sponsored slot optimally assigns the offers it considers best to that slot while deliberately degrading the organic list. It does so only up to the point where end users still find it searching on the platform worthwhile. Degradation is therefore not merely a by-product of monetization but the instrument that makes the sponsored slot worth paying for. The striking implication is hence an inversion of the intuitive hierarchy: The paid slot can be the most informative element on the page, precisely because the platform has degraded the free alternative. Again, the alignment condition depends on the platform's revenue mix. A platform that monetizes only the sponsored slot cares about what business users will pay for it, not about whether the recommended offer suits the end user; one that also earns commissions on completed sales internalizes conversion, and hence fit, when allocating the slot—end users may then benefit from sponsored positions.[56]

The third mechanism concerns the interdependence between selling prominence and the fee structure—the platform's choice of how to monetize steering. The platform may assign prominence by algorithm while taxing transactions, or it may auction prominence directly to business users. These instruments are partial substitutes for extracting the value created by steering. Combining transaction fees with auctions for sponsored prominence gives two revenue streams to the platform. Auctions therefore tend to be associated with lower transaction fees, while algorithmic steering supports monetization through transaction fees. Auctions are attractive where the platform lacks information about demand and cost conditions that bids can elicit; algorithmic steering is attractive where business users have market power or end users differ in their susceptibility to steering.[57] The choice of instrument also affects who ultimately bears the cost: Even where the steering algorithm selects the cheapest offer, end-user prices may be high, because fees and auction payments are passed through. The substitutability of the instruments also foreshadows the anticircumvention analysis: A gatekeeper constrained in its organic ranking may shift to sponsored modules.[58]

The foregoing analysis also illustrates the limits of disclosure obligations. A "sponsored" label may counter the deception risk because end users are told that placement was paid for. It does not, however, eliminate the competitive concerns identified above. A labeled slot still occupies the most prominent position and consumes scarce attention, while disclosure does nothing to restore the informativeness of the surrounding organic list. The U.S. Federal Trade Commission's pending complaint against Amazon illustrates this degradation pattern. It alleges an increase in paid placements and less-relevant sponsored content that, according to the agency, degrades the organic shopping experience while making paid advertising more important for sellers seeking visibility.[59] Concerns are therefore greatest where end users continue to interpret sponsored placement as independent advice while the surrounding organic

---

[56] Janssen et al. (2026), who show that a platform with rich consumer data assigns the products it deems best to the sponsored position and obfuscates organic positions subject to end users' participation constraint; the welfare effect of sponsored positions varies with the platform's objectives, and where the platform earns revenue from sponsored positions together with sales commissions, consumers benefit from sponsored positions.
[57] Bar-Isaac & Shelegia (2026).
[58] See below Section IV.C.4.b).
[59] *Federal Trade Commission 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.), in particular paras. 95–98, 231–236; allegations not adjudicated. See also Scott Morton (2023), 9 fn. 8 (discussing FTC Complaint paras. 5 and 236 and the alleged degradation of search quality associated with increased advertising) and 10 (discussing the displacement of highly ranked organic results by advertised products). For a critical assessment of the FTC's advertising-degradation theory, see Ford (2024).

recommendations have been degraded. In that situation, the sponsored slot inherits precisely the credibility that belongs to the independent-advice benchmark developed in category 1.

Category 6. Advertising Inventory Generation

This objective seeks to capture value by generating end users' attention as an advertising inventory that can be sold to advertisers. The recommender's paying customers are therefore advertisers. The feed bundles organic recommendations with advertising, and the end users pay not with money but with their attention.[60] The defining feature of this category is the revenue base. The platform monetizes the advertising inventory created by end users' attention rather than the recommendation slot itself (category 5) or future retention (category 8).[61] Advertising is the implicit price that end users pay for the ad-financed bundle. That is not necessarily excessive: The platform trades off advertising load against the attractiveness of the bundle, and whether equilibrium advertising is too high or too low relative to the social optimum depends on how much users dislike it relative to the surplus generated by advertising.[62] Market structure also matters. A platform with exclusive access to a user's attention internalizes the scarcity of that attention and therefore avoids overexploiting it. Competing platforms, by contrast, treat users' attention as a common resource. Competition may therefore increase advertising intensity and clutter.[63]

This objective shapes recommendation design in two respects. The value of the feed to the platform equals the number of advertising impressions it generates multiplied by the price that those impressions command. Regarding volume, the platform benefits from feeds designed so that they favor content that increases the amount of advertising inventory generated from the user's current attention. Longer videos that accommodate more advertising breaks provide a particularly clear example. Regarding advertising price, the platform benefits from feeds that favor content that attracts high-paying contextual advertising and therefore yields more valuable impressions. On both counts, organic recommendations become the vehicle through which advertising inventory is created and becomes more valuable. Recommendation design therefore favors content for its advertising value rather than for its value to the end user.

Category 7. Preference Learning

This objective seeks to capture value by learning valuable information about users, products, sellers, creators, or demand. Recommendations elicit clicks, comparisons, saves, purchases, rejections, attention, timing, and willingness-to-pay signals. Part of the value is realized immediately. Better data improves targeting and therefore current advertising revenue, for example by raising advertisers' willingness to bid for sponsored slots. Another part of the value is realized only over time. It is the

---

[60] Peitz (2024) defines attention intermediaries as multisided platforms that attract the attention of viewers and allocate it to attention seekers such as content providers or advertisers, focusing on intermediaries that monetize on the attention-seeker rather than the viewer side.

[61] What separates this category from engagement and retention (category 8) is when the attention is monetized: here immediately, there over the user's future engagement with the platform. The two objectives can conflict: Loading the current visit with monetizable impressions degrades the experience and can cost tomorrow's visits, while a retention-oriented feed may deliberately run lean on advertising today to protect future attention.

[62] Anderson & Coate (2005), the seminal welfare analysis of ad-financed media, in which advertising levels can be too high or too low relative to the social optimum. Anderson & Peitz (2026) on ad load choice and end user time on board.

[63] Anderson & de Palma (2012) on competition for attention when attention is scarce and message congestion arises; Anderson & Peitz (2023) on ad clutter and competition.

forward-looking value that gives the category its distinctive experimentation logic.[64] The experimentation logic also explains where the wedge between platform payoff and end-user value arises. A recommendation may be chosen not because it is the best current match but because the response is informative. The end user exposed to the experimental recommendation bears the cost of learning, while its benefits may accrue to the platform, advertisers, and other users.

The outcome may be relatively benign where learning feeds back into better recommendations for the same users. The mechanism is more problematic where the data serves uses that may extract rather than create end-user value and that connect this category to end-user-side discrimination (category 4). Because recommendations steer the data-generation process itself, a feedback loop can arise: Visibility generates data; richer data improves recommendations; data-enriched recommendations attract more usage; and greater usage generates still more data.[65] This self-reinforcing learning advantage is one of the concerns addressed by the DMA's data-access provisions.[66]

C. Objectives Aimed at Future Value Creation (Categories 8–12)

The second group comprises objectives that use recommendation design to create or improve future opportunities for value capture. They include increasing engagement and retention (category 8), shaping perceptions (9), increasing ecosystem dependence (category 11), and preserving the platform's gateway position (category 12). The last two are closely related but distinct: category 11 operates on users' accumulated ties to the ecosystem, category 12 on the viability and attractiveness of routes around the platform. Operational demand steering (category 10) also belongs to this group insofar as recommendation design is used to improve not only today's operations but tomorrow's market environment of the platform.

These objectives are often more difficult to infer from observed recommendation practice than those aimed at immediate value capture as comprised in the first group. The benefits to the platform are delayed and may arise only indirectly. A demotion designed to weaken a potential rival (category 12), for example, need not increase the platform's current revenue at all; its value lies in preserving future earning opportunities. Recommendation design thus becomes an instrument not merely of current monetization but of broader platform strategy and governance.

Category 8. Engagement and Retention

This objective seeks to increase future opportunities for value by keeping end users engaged with the platform over time. Recommendation design is thus used to increase retention, reduce churn, and make repeated platform use part of users' routines. The defining feature is delayed monetization. Recommendation design may increase current engagement but the value sought by the platform lies in the resulting future usage that in turn is expected to generate future advertising fees, subscription

[64] Bergemann & Bonatti (2019); Hagiu & Wright (2023); Che & Hörner (2018), who show that even a recommender maximizing user welfare optimally over-recommends—pooling exploration with informed advice—to induce the experimentation from which future users learn.
[65] Hagiu & Wright (2023) on data-enabled learning as a source of competitive advantage; Belleflamme & Peitz (2021). As a by-product, biased visibility may also weaken rivals' learning opportunities. Where this becomes the aim rather than a by-product, the conduct falls under preserving the gateway (category 12).
[66] See below Section IV.C.2.

income, or data value.[67] Personalized playlists, infinite feeds, notifications, and "because you watched/bought" recommendations illustrate the mechanism. They may prolong users' current engagement but, more importantly for this category, they also encourage repeated and sustained use of the platform over time.

In this category, the source of potential bias and decreased consumer welfare differs from that in the preceding categories. There, the concern generally arises because recommendation design pursues a platform objective that may diverge from what end users would choose. Engagement optimization creates a different challenge: Recommendations may faithfully reflect users' observed choices and still fail to advance what users themselves ultimately value because in-the-moment behavior is an unreliable guide to what users truly value. Time spent is not value received: Users are often only imperfectly aware of their self-control problems. Field-experimental evidence attributes around 30% of social media use to self-control problems.[68] Observed end-user engagement with a platform is therefore an imperfect proxy for consumer welfare.

The implication is striking. Where behavior and reflective preferences diverge, even a platform that genuinely seeks to maximize user value, but does so only by observing engagement, may optimize for the users' impulsive selves rather than their reflective selves.[69]

The same concern becomes even stronger where the divergence between engagement and end-user welfare is deliberately increased through service design. Platforms choose how habit-forming their services are. Where attention is scarce, addictiveness is the instrument by which each platform pulls attention away from rivals, and this business-stealing incentive leads competing platforms to choose greater addictiveness than a monopolist would.[70] Retention-oriented recommendation design is therefore benign where repeated usage reflects sustained end-user value. It becomes harmful where engagement is manufactured through habit formation and exploited self-control problems, so that time on the platform no longer reveals value to the user but instead reflects addictiveness and other behavioral distortions.

Category 9. Perception Shaping

For this objective, the platform uses recommendation design to shape end users' beliefs, associations, tastes, perception of quality, and salience regarding the platform's own brand and products, its ecosystem, or third-party products and services the platform monetizes. Recommendation thus becomes an instrument of branding and demand creation. Rather than merely steering choices among current alternatives, it influences how these alternatives are perceived. The payoff lies in stronger future

[67] The boundary to advertising inventory generation (category 6) lies in the timing of monetization; see the discussion there.
[68] Allcott, Gentzkow, & Song (2022), based on a randomized field experiment on smartphone use.
[69] Kleinberg, Mullainathan, & Raghavan (2024), who model media consumption by users with inconsistent preferences—an impulsive process governs in-the-moment engagement; a reflective process governs what users actually value—and show that the problem transcends platform incentives: Even a platform maximizing user utility on the basis of observed engagement inherits the bias, and whether higher engagement raises or lowers user utility depends on the content changes that produce it. See also Fudenberg & Levine (2006) for a dual-self model.
[70] Ichihashi & Kim (2023). In their model, addictiveness is a design choice: A more addictive service yields lower utility from participation but a higher marginal utility of allocating attention to it, so addictiveness holds the attention that the platform monetizes at the expense of the value that users derive.

demand rather than today's transaction.[71] Examples include editors' picks, curated collections, lifestyle rankings, and product badges.

This category differs from end-user-side discrimination (category 4) in the object of recommendation design. There, the platform exploits users' given heterogeneity in willingness to pay. Here, recommendation design seeks to shape willingness to pay itself by influencing the beliefs, tastes, and perceptions from which future demand emerges. The mechanism operates through the composition of what is displayed to users. Because valuations are context-dependent, and end users tend to overweight the attributes that the choice environment makes salient, the platform can engineer perceived quality or perceived bargains by deciding which options appear together and which attributes stand out.[72]

Curation devices borrow the recommender's credibility. An editor's pick, curated collection, or product badge reads as a certification and is informative only insofar as the underlying selection tracks genuine quality. Perception shaping is therefore beneficial to end users where curation reliably signals quality and directs attention to the attributes that truly matter. It becomes detrimental where engineered salience inflates perceived quality or value beyond what end users experience after purchase. The *Coupang* case discussed under category 3 illustrates this mechanism. Manipulated ratings and review counts may shape consumers' perceptions of quality in order to promote the platform's own products.

Category 10. Operational Demand Steering

This objective seeks to improve the platform's efficiency by using recommendation design to manage capacity, inventory, logistics, congestion, timing, and availability constraints. Rather than being guided solely by user-match quality, recommendation design is deployed to align demand with operational conditions. In doing so, it seeks to reduce operational costs, clear inventory, smooth demand, or balance supply and demand. Such recommendations may depart from a pure relevance benchmark while nevertheless increasing overall efficiency. Examples include recommendations for substitute items for out-of-stock products, off-peak travel options, sellers offering faster delivery, currently available drivers, and products with a lower probability of return.[73]

The KFTC's *Kakao Mobility* case illustrates both the practical relevance of this category and the evidentiary challenges it raises. The Korean competition authority found that the dispatch algorithm favored the ride-hailing platform's own franchise taxis by giving them preferential treatment within a predefined estimated time of arrival range and by applying acceptance-rate criteria that disadvantaged nonfranchised drivers. According to the authority, this allowed the platform to leverage its dominance in ride hailing into the taxi-franchise market: Higher ride volumes and revenues increased drivers' incentives to join the franchise, thereby strengthening network effects and lock-in.[74] The Seoul High Court, however, overturned the decision, finding that the differential treatment did not constitute

[71] Where the shaped perception serves a third party's persuasive aims or the platform's own convictions rather than a commercial objective, the conduct falls under categories 13 or 14.

[72] Bordalo, Gennaioli, & Shleifer (2013), in whose salience framework the attributes consumers overweight (i.e., quality or price) depend on the composition of the choice set, so that the same good can be made to look like a bargain or a rip-off by varying its context.

[73] Cui & Shin (2018); Cui, Zhang, & Bassamboo (2019); von Wangenheim (2025); Caldentey & Xie (2026).

[74] Korea Fair Trade Commission (KFTC), 13.6.2023, Decision No. 2023-093, Case No. 2021Sigam0760, *Kakao Mobility*. See OECD, *Monopolisation, Moat Building and Entrenchment Strategies – Note by Korea*, DAF/COMP/WP3/WD(2024)6, paras. 9–15; Choi & Fuchikawa (2024), 505–506.

unlawful discrimination and, in any event, that the KFTC had failed to sufficiently establish anticompetitive effects in the taxi-franchise market.[75]

Dispatch and allocation algorithms illustrate why recommendation design can only be understood and classified by reference to the platform's objective. The same design feature may serve operational demand steering (an efficiency rationale based on expected pick-up times) or the promotion of own services and ecosystem dependence (a strategic rationale). As a reminder, the taxonomy classifies objectives rather than design levers. This distinction may also be legally significant. While both objectives may deviate from the benchmark of providing independent advice that maximizes end-user benefit, different considerations may be required for their legal assessment.

Category 11. Ecosystem Dependence

This objective seeks to increase the dependence of end users, sellers, creators, developers, or complementors on the platform ecosystem. Recommendation design is used to steer users toward products, services, or complements that deepen ecosystem ties, so that the cost of leaving rises. The concern is that dependence on the platform ecosystem grows, such that tomorrow a superior alternative may fail to be adopted not because it is disadvantaged but because leaving has become too costly. Throughout, the design works on the user's accumulated ties to the ecosystem so that the cost of leaving rises even where each individual recommendation is sound. Examples include recommendations for the platform's own payment service, logistics network, app, operating-system features, or subscription bundle, as well as exclusive shows, games, or superstar creators.

Recommendation design can build ecosystem dependence through three distinct mechanisms. The first mechanism is *ecosystem-specific investment*. Recommendations steer end or business users toward proprietary services that require end users or business users to invest resources that have little value outside the ecosystem, such as logistics integration, purchased content libraries, learned workflows, and accumulated reputations. As a result, leaving the ecosystem means writing off those investments. Personalization creates a functionally similar form of lock-in. Where end users' own usage data is used to tailor a service, the platform's services become progressively better suited to those end users than competing services. Recommendation-steered usage can therefore increase switching costs simply by generating additional personalized data.[76]

The second mechanism is ecosystem *envelopment*. The platform expands into adjacent services by leveraging its existing user relationships and shared data, and prominence generated by the recommendation policy acts as a virtual bundling instrument, steering the installed user base toward those adjacent services. In this way, users become tied to an increasing number of complementary services, so that the cost of exiting the ecosystem rises with every additional service adopted.[77]

[75] Seoul High Court, 22.5.2025, Case No. 2023Nu50891, *Kakao Mobility v Korea Fair Trade Commission*; appeal pending before the Supreme Court of Korea, Case No. 2025Du34036.

[76] Hagiu & Wright (2023), who distinguish across-user learning (each user's data improve the service for all, cited under category 7) from within-user learning, where a user's own data improve the service for that user—the latter operating like a switching cost.

[77] Eisenmann, Parker, & Van Alstyne (2011) on platform envelopment: entry into an adjacent platform market by bundling the target's functionality with the attacker's, leveraging overlapping user relationships and shared components.

The third mechanism operates through *exclusive content*. Recommending exclusive superstar creators or sellers, shows, or games gives end users an additional reason to join or remain within the ecosystem. Cross-group network effects amplify this dynamic: Exclusivity attracts end users, end users attract complementors, and each group increases the value of the ecosystem for the other.[78]

The welfare assessment of ecosystem-favoring recommendation design turns importantly on whether the resulting dependence primarily reflects superior user value or strategically created switching costs. In the former case, the platform steers users toward ecosystem options that genuinely serve them through seamlessness or quality, with dependence emerging as a by-product of the value created. In the latter, it steers them toward options chosen for the switching costs they generate, allowing the platform to monetize the resulting dependence later through fees, prices, or user attention.[79]

The Amazon logistics proceedings illustrate ecosystem-specific investment and switching costs. In Italy, the competition authority found that Amazon tied access to Prime-related sales and visibility advantages to the use of its Fulfillment by Amazon (FBA) service. Those advantages included the Prime label and an increased likelihood of selection for the buy box, both of which were central to sellers' commercial success. According to the authority, Amazon thereby leveraged its marketplace power into logistics while increasing sellers' ecosystem-specific commitments: Sellers using FBA faced higher costs of multihoming because maintaining a comparable presence on rival marketplaces would require duplicating warehousing arrangements.[80] The Italian remedy sought to unwind this mechanism by requiring Amazon to make sales and visibility benefits available to third-party sellers meeting fair and nondiscriminatory fulfillment standards, irrespective of whether they used FBA. The subsequent EU[81] and UK[82] proceedings addressed closely related mechanisms through commitments aimed at reducing the dependence of buy box prominence and Prime participation on Amazon's own logistics services. In addition, Amazon committed not to use nonpublic seller data for its own retail decisions, to apply nondiscriminatory criteria in buy box selection, and to allow Prime sellers greater freedom to use and negotiate with independent delivery providers.

---

[78] Carroni, Madio, & Shekhar (2024), who analyze when a platform secures exclusivity over a superstar and how cross-group network effects amplify the exclusive's effect on participation.

[79] The boundary to engagement and retention (category 8) runs along the mechanism of the attachment: category 8 keeps users returning through the appeal of the experience; here the platform raises the cost of leaving.

[80] Autorità Garante della Concorrenza e del Mercato (AGCM), 30.11.2021, Case A528 – *FBA Amazon*, Decision No. 29925, in particular paras. 253–254, 698, 702, 737, 826–831; see, for an English-language account of the decision, AGCM, 9.12.2021, Press Release, "Italian Competition Authority: Amazon fined over €1,128 billion for abusing its dominant position." See also OECD, *Annual Report on Competition Policy Developments in Italy – 2021*, DAF/COMP/AR(2022)19, paras. 55–64; Colangelo (2023) 549–550; Falce & Faraone (2023) 48–51. For a critical assessment, see Andreoli-Versbach & Gans (2024), 472–488. The infringement finding was subsequently upheld by the TAR Lazio, Sez. I, 2.9.2025, Judgment No. 15919/2025, Case No. 1364/2022, while the decision was partially annulled as regards the calculation of the fine. An appeal is pending before the Consiglio di Stato, which has referred to the Italian Constitutional Court questions concerning the time limit for the AGCM to initiate antitrust proceedings; Consiglio di Stato, Sez. VI, 30.6.2026, Order No. 5167/2026, Case No. 9396/2025.

[81] European Commission, 20.12.2022, Cases AT.40462 – *Amazon Marketplace* and AT.40703 – *Amazon Buy Box*, Decision C(2022) 9442 final, in particular paras. 232–236; Summary of Commission Decision, OJ 2023 C 87/7.

[82] Competition and Markets Authority (CMA), 3.11.2023, Case No. 51184, *Investigation into Amazon's Marketplace*, Decision to Accept Binding Commitments under the Competition Act 1998 from Amazon in Relation to Conduct on its UK Online Marketplace; monitoring trustee approved 5.2.2024.

The same ecosystem-building dynamic can also be achieved through defaults and preinstallation. The Japan Fair Trade Commission (JFTC) ordered Google to cease licensing and revenue-share terms that required preinstallation and prominent placement of Google Search and Chrome on Android devices and, under certain revenue-sharing arrangements, restricted the implementation and promotion of rival search services.[83] Insofar as the same terms excluded rival search services, the case also implies gateway preservation (category 12), to which we turn next.

Category 12. Preserving the Gateway

This objective seeks to preserve the platform's role as the primary gateway between end users and business users. Recommendation design is used to discourage users from switching to or even considering alternative routes through which rivals or complementors could reach them. As with ecosystem dependence (category 11), the distinctive concern is dynamic: Today's recommendations are made to reduce the viability of competing access routes and thereby preserve the platform's gatekeeper position.

The two categories differ in the object of the design. Category 11 works on the user's ties to the ecosystem, increasing what would be lost on leaving. Category 12, by contrast, operates on the alternative routes themselves, reducing what end or business users stand to gain from taking them. In the vocabulary of Section II, gateway preservation is the objective that most directly protects the effective-access bottleneck from which intermediation power derives. Examples include demoting direct-channel offers, discouraging off-platform payments, and steering users away from rival app stores, booking channels, or specialized search services. The alternative routes at issue are of three kinds: rival intermediaries, business users' own direct channels, and off-platform completion of transactions initiated on the platform.

The first mechanism for achieving this objective is *weakening rival gateways*. By reducing their visibility, recommendation design deprives rivals and complementors of traffic. This, in turn, limits the scale and data on which discovery quality depends, preventing specialized search services, rival app stores, or aggregators from ever reaching the quality at which users would switch routes. The logic is not confined to discovery: Steering fulfillment toward the platform's own logistics network denies rival carriers the delivery density on which speed and cost depend, so that they too may never reach the quality at which sellers would switch. The same objective may also be pursued by acquiring the very tools through which users compare platforms and steering users through them.[84] The distinction turns on the effect for which the design is adopted: the additional ecosystem tie in category 11 or the weakening of the adjacent rival gateways in category 12.

[83] Japan Fair Trade Commission (JFTC), 15.4.2025, Cease and Desist Order No. 5 of 2025, *Google LLC*; see JFTC, *JFTC Issues a Cease and Desist Order to Google LLC*, 15.4.2025 (provisional English translation); Lee (2025).

[84] Carlton & Waldman (2002), whose logic—foreclosing a complementor today to prevent it from evolving into a rival tomorrow—can be translated from tying to steering. Motta (2023) on exclusionary theories of harm for self-preferencing, running through rivals' loss of scale and data. Athey & Scott Morton (2022) on "platform annexation": a platform acquires or controls a complementary tool through which users compare and move across platforms (e.g., search engines, wallets, ad servers) and operates it to steer users toward itself—through prominence and defaults within the tool, or by making rivals' options costlier to reach through degraded access. Initially small diversions of traffic then compound as rivals lose the scale on which their quality depends.

The second mechanism is *penalizing off-platform differentiation*. Visibility penalties imposed on business users offering better terms elsewhere make multihoming and direct channels costly and less attractive. Recommendation design thus protects the platform's fee structure and prevents business users' own channels from developing into viable routes around the platform. The conversion-rate mechanism described under category 2 supplies the design lever; here it is deployed for a different objective, namely preserving the gateway.

The third mechanism is *preventing disintermediation*. Recommendation design steers users toward completing transactions on the platform, for example through payment, communication, or contracting. By keeping both the end-user relationship and the transaction data within the platform, recommendation design preserves the informational and commercial advantages that make the gateway position durable.[85] Steering toward a proprietary payment or communication service may alternatively or in addition reflect ecosystem dependence (category 11), own-offer monetization (category 3), or the value of the transaction data it yields (category 7). The gateway objective is implicated only where the value of the design lies in preventing alternative routes from becoming viable.

The gateway-preservation objective is most readily observed in visibility penalties imposed on off-platform activity. Hotel booking platforms illustrate gateway preservation through ranking penalties attached to off-platform pricing. It has been documented that hotels offering lower prices outside the platform tend to receive worse positions in the organic ranking. Such ranking penalties discourage direct booking and thereby preserve the platform's role as the gateway between hotels and consumers. We return to this pattern in the discussion of the DMA's anticircumvention provisions in Section IV.C.4.a).[86] The Amazon proceedings in Germany illustrate a similar strategy in the context of marketplace pricing. The Bundeskartellamt, the German competition authority, found that Amazon employed algorithmic price-control mechanisms that deactivated offers or excluded them from the Featured Offer (buy box) when prices exceeded internally determined thresholds or benchmarks derived from prices charged by external online retailers. According to the authority, these visibility penalties exerted substantial pressure on marketplace sellers to adjust their prices to Amazon's benchmark, with exclusion from the buy box carrying major commercial consequences. The authority further considered that this practice could weaken price competition beyond Amazon's marketplace by discouraging independent pricing initiatives by other online retailers.[87] The U.S. FTC's complaint against Amazon shows how multiple visibility mechanisms may operate together in support of the same underlying objective. The complaint alleges that Amazon maintained its monopoly through buy box disqualification, search demotion, price hiding, exclusion from advertising and recommendation widgets, and Prime/FBA conditioning. According to the agency, the visibility penalties and fulfillment conditioning reinforced one another, deterring sellers from supporting alternative sales channels and limiting rivals' ability to gain the scale necessary to become viable alternatives.[88] Prime and FBA conditioning also featured in the European proceedings discussed under category 11; the difference lies

---

[85] Hagiu & Wright (2024), already cited under category 2, on platform responses to marketplace leakage.
[86] Hunold, Kesler, & Laitenberger (2020); Franck & Peitz (2024), section 4.3.2.
[87] Bundeskartellamt, 4.2.2026, Case No. B2-73/20, *Amazon Marketplace—Price-Control Mechanisms*, in particular paras. 365–378, 795–805, 822–835. Amazon has challenged the decision before the German Federal Court of Justice (Bundesgerichtshof); proceedings are pending.
[88] *Federal Trade Commission 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.), in particular paras. 279, 284–287, 360–365, 412–416; allegations not adjudicated. For discussion, see Scott Morton (2023), 3–8.

in the objective asserted, not in the conduct. Exclusive content, the third mechanism under ecosystem dependence (category 11), may equally be deployed to preserve the gateway. In South Korea, Google allegedly used preferential visibility on Google Play to induce game developers to withhold important new titles from the rival app store One Store. According to the KFTC, Google offered prominent featuring and other promotional support on Google Play in exchange for exclusive releases.[89]

D. Influence, Welfare, and Integrity (Categories 13–16)

The third group comprises platform objectives for which recommendation design serves purposes other than immediate or strategic value capture. Rather, the platform uses recommendations to serve persuasive intent pursued by third parties or itself, advance public-interest objectives, or protect the platform's integrity and safety. The platform's underlying motivations vary. Some objectives remain commercially driven, such as the platform's neutrality toward third-party persuasion (category 13) or risk management (category 16); others reflect the platform's own normative convictions (category 14); and public-interest objectives (category 15) may arise from the desire to comply with legal obligations or with genuinely internalized societal goals. Recommendation design may therefore depart from the independent-advice benchmark without seeking either immediate monetization or the creation of future opportunities for value capture that go beyond the platform's economic viability as such. Such departures do not necessarily amount to undue favoritism. The bias arising from the departure from the independent-advice benchmark may represent manipulation in one setting, user protection or public-interest promotion in another, and a necessary safeguard of platform integrity in a third.

Category 13. Third-Party Persuasion

This objective seeks to preserve engagement, content supply, and advertising revenue by accommodating, rather than policing, third parties' attempts to persuade end users for noncommercial purposes via the recommender. Thus, unlike the following category 14, the persuasive purpose originates with a third party that seeks fame, popularity, social influence, or ideological reach. The platform's own motivation is profit: tolerating such behavior may be more profitable than policing it because moderation is costly and the resulting content often generates particularly high engagement. What distinguishes this category is that recommendation design amplifies persuasion initiated by third parties rather than by the platform itself. The platform provides a means of showcasing particular content through rankings, popularity signals, badges, snippets, or endorsement-like presentation. The third party creates prominence by manipulating organic signals, while the platform chooses to tolerate rather than correct such manipulation. Examples include coordinated amplification of ideological content, astroturfed popularity, and engagement-bait content whose artificially inflated prominence the platform leaves uncorrected.[90] The resulting recommendation bias stems from omission rather than

[89] Korea Fair Trade Commission (KFTC), 20.7.2023, Decision No. 2023-103, Case No. 2020Jigam1764, *Google App Store*, paras. 55–76; 96–214; for an English-language account of the decision, KFTC, 11.4.2023, Press Release, "KFTC Sanctions Google for Anticompetitive Conduct in App Market Blocking Developers from Releasing Games on Competing Platform." See also OECD, *Monopolisation, Moat Building and Entrenchment Strategies – Note by Korea*, DAF/COMP/WP3/WD(2024)6, paras. 16–24; Lee (2024).

[90] Allcott & Gentzkow (2017) on the supply of and demand for fabricated news in the 2016 election, spread through social media sharing and ranking; Vosoughi, Roy, & Aral (2018), who document in Twitter data that false news diffuses farther and faster than true news—content that performs well organically, which is what makes accommodation profitable.

intervention. It arises where the platform knowingly allows manipulated organic signals to influence recommendation outcomes because doing so is more profitable than effective moderation.

Category 14. Ideological and Social Influence

This objective seeks to use recommendation design to induce or reinforce a platform-preferred view, attitude, identity, social norm, or behavior. Unlike in the previous category 13, the persuasive intent originates with the platform or curator itself. Recommendation design may serve this objective by amplifying selected content through prominence and social reinforcement and contribute to identity formation. In pursuing these aims, the platform may depart from user-match quality and sacrifice engagement or revenue.

Field-experimental evidence demonstrates that algorithmic recommendation is capable of producing such effects. Exposure to an algorithmic feed can durably shift political attitudes, in part by changing whom users follow and thereby reshaping their future information environment.[91] Evidence of such effects does not, however, establish that ideological influence was the platform's objective. An observed political tilt may instead result from engagement optimization (category 8), making this a particularly acute instance of the inference problem discussed in subsequent Section III.E. The category differs from public-interest steering (category 15) because the views or social norms promoted reflect the platform's particular convictions rather than broadly shared or externally anchored public-interest standards. Examples include the amplification of political content to promote a preferred viewpoint, owner-driven ideological curation of feeds, and recommendation designed to foster particular social norms.

Category 15. Welfare and Public Interest

This objective seeks to use recommendation design to improve individual or social outcomes, such as knowledge, health, sustainability, fairness, diversity, or civic capacity. The platform may pursue this objective for various motives: to comply with existing regulation, to preempt future regulation, or because it has internalized societal goals, possibly reflecting the preferences of shareholders or employees.

The category differs structurally from the preceding one. There, departures from individual match value or user benefit serve objectives centered on the platform's own interests or preferences. Here, recommendation design takes account of welfare effects beyond those captured by the individual match value or user benefit. Those broader effects include externalities such as the civic value of a diverse information environment or the health and sustainability consequences of consumption choices. Recommendations may accordingly depart from what the individual user would choose precisely

[91] Gauthier, Hodler, Widmer, & Zhuravskaya (2026), whose field experiment on X finds that seven weeks of exposure to the algorithmic feed—which favored conservative and activist posts while demoting traditional news outlets—durably shifted users' political attitudes, the persistence arising because the algorithm changed which accounts users followed; as the authors note, the algorithm prioritized right-leaning content already under the platform's previous ownership, cautioning against attributing the tilt to owner conviction. Allcott, Braghieri, Eichmeyer, & Gentzkow (2020), whose randomized Facebook-deactivation experiment finds reduced political polarization and modestly higher subjective well-being but also reduced factual news knowledge. Hosseinmardi et al. (2021), who find that consumption of radical content on YouTube is driven more by user demand than by algorithmic amplification—a further caution against reading exposure patterns as a platform objective.

because doing so serves the general public interest.[92] Whether such a departure counts as bias at all depends on the benchmark against which recommendations are assessed (Section III.A).

The significance of this category extends beyond consumer choice. Rankings allocate visibility and prominence not only among products but also among sellers, creators, workers, candidates, publishers, and viewpoints. Recommendation design can therefore determine how opportunities are allocated and shape the quality and diversity of the information environment. Thus, recommendation design may work as an instrument of opportunity allocation and epistemic quality. Examples include recommendations that guide users along learning paths or toward healthier choices, exposure-diversity rules that broaden the range of information users encounter, measures that ensure the findability of public-service media content, and recommender designs that facilitate long-tail discovery.

Google, for example, claims that its search algorithm ranks results according to their expected usefulness to the individual user, considering not only the relevance of the content but also the "experience, expertise, authoritativeness, and trust" ("E–E–A–T") of its source.[93] According to the platform's publications, the latter considerations carry particular weight for sensitive topics "that have a high risk of harm because content about these topics could significantly impact the health, financial stability, or safety of people, or the welfare or well-being of society" ("Your Money or Your Life" topics). These include medical and health-related content but also "topics that could negatively impact groups of people, issues of public interest, trust in public institutions, election and voting information, and any other informational topics about government, civics or society that impacts people's lives, government information, elections, and matters affecting societal trust."[94] The criteria for quality raters are set out in Google's Search Quality Rater Guidelines. Their ratings do not directly determine search rankings; rather, they provide feedback used to evaluate and improve Google's ranking systems.[95]

A further instructive example for this category emerged from the *netdoctor v Google* case, a successful private antitrust enforcement against a cooperation between Google and the German Federal Ministry of Health.[96] They had agreed that, for Google searches relating to certain diseases, content from a government-operated national health portal would be displayed exclusively in prominently placed knowledge panels, with links directing users to the portal. The proceedings revealed that this arrangement built on a prior general decision by Google to display prominent health-information boxes as it also did, for example, in the United States and the United Kingdom.[97] Google maintained that the

[92] Even a neutral design in the sense of category 1 may fail a public-interest benchmark: Recommendations biased toward popular options reduce the diversity of what users consume (Fleder & Hosanagar, 2009). Serving the objective therefore requires deliberate design—for instance, adopting exposure diversity as a design principle for the recommender (Helberger, Karppinen, & D'Acunto, 2018).
[93] Google (2025), section 3.4.
[94] Google (2025), section 2.3.
[95] Google (2025), section 0.1 ("No single rating can directly impact how a particular webpage, website, or result appears in Google Search, nor can it cause specific webpages, websites, or results to move up or down on the search results page … Instead, ratings are used to measure how effectively search engines are working to deliver helpful content to people around the world. Ratings are also used to improve search engines by providing examples of helpful and unhelpful results for different searches").
[96] Regional Court of Munich I, 10 February 2021, Case 37 O 17721/20, *netdoctor/Google*, Juris. The decision became final when Google withdrew its initial appeal. See Persch (2021) for an analysis of the key antitrust aspects of the decision.
[97] Regional Court of Munich I (n 96), Juris, paras. 10 and 44.

purpose was to provide users with reliable health information,[98] something it considered could not be ensured by the search algorithm alone, notwithstanding the E–E–A–T-based assessment of sources described above. Google further submitted that the sources used for its knowledge panels were selected with the assistance of medical experts, who assessed the information for reliability and comprehensibility; Google emphasized that the selected information was evidence-based, professionally reviewed, and free from commercial and political interests.[99]

The case illustrates four aspects that matter for our taxonomy. First, Google's lever—preferential visibility for hand-picked public-health information—involved a deliberate departure from ordinary relevance-based ranking, accompanied by a substantial loss of visibility for competing health-information portals. Second, the use of this lever does not by itself reveal whether there was any bias relative to category 1: Were users better served by Google's curated health information than by the results that its ordinary search algorithm would otherwise have selected? Google's position was essentially that they were but that depends on empirical findings and the specifics of the normative benchmark. Third, Google also invoked a broader public interest in access to reliable health information.[100] Therefore, if we assume that the curated recommendations departed from what individual search users themselves would have preferred to see and, therefore, departs from the independent-advice benchmark, the case would provide an illustrative example of category 15: recommendation design departing from individual match value in pursuit of a broader public-interest objective. Fourth, Google might also have used the knowledge tools to keep end users on the platform for longer, thereby pursuing an engagement-and-retention objective within the meaning of category 8. Where exposure to certain material instead raises concerns about platform integrity or safety, recommendation design falls within category 16, which closes our taxonomy.

Category 16. Platform Integrity and Safety

This objective consists in safeguarding the integrity and safety of the platform by reducing users' exposure to unlawful, unsafe, deceptive, or otherwise harmful content, products, or actors. The platform may pursue this objective to protect end users or the integrity of its service but also to reduce legal and regulatory risk or protect its political and reputational standing. This objective therefore sits at the boundary with the second group of objectives (Section III.C).

Recommendation design may serve this objective by favoring verified, authoritative, lawful, or safe options and demoting or even delisting content, products, or actors associated with greater risk, including where the risk results from commercially motivated manipulation such as fake reviews or counterfeit products.[101] Apple's App Store provides a concrete illustration: Apple has expressly justified

---

[98] Regional Court of Munich I (n 96), Juris, paras. 8 and 46.
[99] Regional Court of Munich I (n 96), Juris, paras. 56–57.
[100] Regional Court of Munich I (n 96), Juris, para. 60.
[101] Category 16 mirrors category 13: There the platform tolerates the manipulation of organic signals; here it polices it. The policing may be a response to a documented problem. Ott, Cardie, & Hancock (2012) estimate the prevalence of deceptive opinion spam in online review communities; Luca & Zervas (2016) find that a substantial share of Yelp reviews is filtered as suspicious, with fraud more likely when reputations are weak and unfavorable fake reviews more likely under stronger competition; Mayzlin, Dover, & Chevalier (2014) show that review manipulation varies with posting requirements and with the incentives of the parties that stand to gain from it. On the theory side, Dellarocas (2006) models sellers' incentives to manipulate internet opinion forums; Resnick & Sami (2007, 2008) show that recommender systems can be designed to resist manipulation, but only at a cost in informativeness.

excluding deceptive apps from search results and App Store charts as necessary to protect the "integrity of app discovery," while its App Review Guidelines provide for the removal of applications presenting safety, fraud, privacy, or other specified risks.[102] Such measures may protect users and the integrity of the service, but they may at the same time reduce the platform's exposure to fines, liability, reputational harm, or stricter regulation.

This illustrates that incentives to pursue the platform-integrity objective may be externally induced. Where the platform faces liability for harm caused by third-party offers, for example, screening and demotion acquire a direct pecuniary rationale.[103] Such recommendation design may depart from a pure relevance benchmark while remaining fully aligned with end-user interests: An option may be less relevant to an individual user yet appropriately receive less prominence because it poses a greater risk of harm. The normative assessment therefore turns on whether the asserted integrity or safety concern is genuine and whether the recommendation response tracks that risk. Integrity-oriented ranking is benign where demotion responds to genuine risk and thereby operates as quality control in the users' interest. It may deserve scrutiny where integrity or safety concerns serve as a pretext for pursuing potentially problematic objectives such as the demotion of rivals (category 12) or steering toward the platform's own services (category 3).

E. Design Levers and Recommendation Objectives: Cross-Cutting Observations

The taxonomy developed above reveals that recommendation design levers cannot be mapped one-to-one onto objectives: The same design lever may serve different objectives, while the same objective may be pursued through different designs. The principal design levers include ranking and filtering; defaults; badges, labels, and snippets; sponsored slots; personalization; notifications and timing; and suppression and demotion. These levers cut across the taxonomy of objectives. Self-preferencing, for example, may serve to monetize the platform's own offers (category 3) or to increase ecosystem dependence (category 11); personalization can improve user fit (category 1) or facilitate surplus extraction (category 4); sponsored slots may monetize visibility (category 5) or facilitate third-party persuasion (category 13); and demotion may protect platform integrity (category 16) or weaken competing gateways (category 12).

Google's integration of AI overviews at the top of the search results page provides a useful illustration. The prominence of the AI-generated answer is plainly an exercise of recommendation design, yet the underlying objective cannot be inferred from prominence alone. First, the design may simply pursue the independent-advice objective in category 1: Synthesizing information directly on the results page may reduce search costs and provide users with a better and faster answer than requiring them to inspect and compare several individual search results.[104] Second, engagement and retention (category 8), ecosystem dependence (category 11), and gateway preservation (category 12) provide plausible

---

[102] Apple, "The App Store stopped over $2.2 billion in fraudulent transactions in 2025" (20 May 2026) (reporting that, "[t]o further protect the integrity of app discovery," Apple blocked deceptive apps from App Store search results and charts), https://www.apple.com/newsroom/2026/05/the-app-store-stopped-over-2-point-2-billion-usd-in-fraudulent-transactions-in-2025/; Apple, App Review Guidelines (June 8, 2026), §§ 1, 2.3.1, 5.1, 5.6.3., https://developer.apple.com/app-store/review/guidelines/.

[103] Hua & Spier (2025) analyze platform liability for harms caused by third-party products and the screening incentives induced by such liability.

[104] This corresponds to Google's own account of the objectives underlying AI Overviews. See Google's Vice President, Product Management, Search, Budaraju (2025).

additional hypotheses. By answering more of the user's query within Google Search itself, AI overviews may increase the amount and duration of interaction with the service and encourage users to use it for increasingly complex information needs (category 8).[105] By placing a synthesized answer above the organic results rather than sending users on, AI overviews may strengthen search as the single point of access to Google's own vertically integrated services (category 11). Their prominence may also reduce users' occasions to turn to alternative AI-based answer services, which, receiving fewer queries, may never reach the quality at which users would switch (category 12). Further objectives are conceivable but require additional assumptions. Own-offer monetization (category 3) would become relevant if the prominence of the AI overviews were intended to steer users toward a proprietary Google service from which Google derives greater value than from referral to third-party content; prominence of Google-generated content alone does not establish that objective. Advertising inventory generation (category 6) is plausible insofar as AI-enhanced search creates additional or more valuable opportunities for advertising, but this requires a link between the design and the generation of monetizable advertising attention. Preference learning (category 7) may likewise matter where interactions with AI overviews generate richer information about users' interests and intentions that Google values for future recommendation, search, or advertising purposes. Again, however, the fact that the design generates such data does not establish that learning from them is the objective for which the design was chosen. The example thus illustrates the inferential gap: The same prominent design feature may be consistent with user-oriented advice, engagement, ecosystem dependence, gateway preservation, monetization, or learning, and identifying the objective requires evidence beyond the observed interface.

1. The Inferential Gap Between Design Levers and Objectives: Counterfactuals, Incentives, and Effects

Because observed design does not uniquely identify the objectives it serves, and the set of possible objectives need not be confined to those distinguished here, isolating the design lever can only be the starting point of an analysis of recommendation power. Identifying the objective requires additional evidence about the platforms' incentives, the relevant counterfactual ranking, and the effects of the observed design. The counterfactual asks which offers would receive prominence if ranking were based on relevance, quality, price, or expected welfare. The incentive analysis asks whether observed prominence is systematically related to commissions, royalties, paid placement, own-offer status, or off-platform leakage. The effects analysis examines how changes in rank affect traffic, conversion, and business-user scale, and ultimately data accumulation, entry incentives, or the viability of direct sales channels.[106] Empirical tests for recommendation bias operationalize this inquiry through counterfactual comparisons. They differ along three dimensions: how prominence is measured, which benchmark is used to construct the counterfactual ranking, and which attributes are treated as legitimately determining rank. The latter two lead directly back to the benchmark question discussed under category 1.[107]

---

[105] According to Google, AI Overviews had by 2025 increased usage of Search by more than 10% for the types of queries for which they were displayed in major markets. See Budaraju (2025). Google also reported that, since the launch of AI overviews, "there's been a profound shift in how people are using Google Search. People are coming to Google to ask more of their questions, including more complex, longer and multimodal questions"; Reid (2025).

[106] For the corresponding evidentiary checklist in abuse cases, see Belleflamme & Peitz (2021); Hagiu, Teh, & Wright (2022).

[107] Reimers & Waldfogel (2023) compare observed ranks with the frontier that maximizes a weighted sum of consumer and seller surplus. Jürgensmeier & Skiera (2025) measure an offer's visibility across an entire platform

2. *Bridging the Inferential Gap Between Design Levers and Objectives*

A taxonomy of objectives helps translate observed recommendation design into hypotheses about the objectives it serves, the underlying platform incentives, and the likely effects, thereby providing a basis for normative assessment. It may thus also inform legal analysis under existing regulatory frameworks. Where, for example, a platform operator possesses sufficient market power to be subject to abuse control under Article 102 TFEU or equivalent provisions, the objective served by a particular recommendation design may be relevant to determining whether the conduct is abusive.

Drawing normative conclusions about when regulatory intervention is warranted, however, presupposes that the relevant objectives can be identified with sufficient confidence in practice. The taxonomy and the illustrations developed above show why this may be difficult in individual cases. Objectives may be closely related, the boundaries between them may be porous, and the same recommendation design may plausibly serve more than one objective. Any attempt to attach normative consequences to particular objectives must therefore address the inferential gap between observable design levers and the objectives they serve.

Bridging this gap may involve both subjective and objective evidence. To identify the objective plausibly associated with a particular design lever, the analysis may consider the platform's own intentions and motivations, as well as the observable effects of the design. Subjective intent is, by its nature, often difficult to verify. In individual cases, however, internal documents or other evidence may reveal the considerations that motivated the platform's design choices and may therefore support the inference that a particular objective was being pursued. Reliance on such evidence is by no means alien to antitrust analysis. Subjective considerations relating to an undertaking's motivation or intent may, for example, form part of the assessment of exclusionary conduct under Article 102 TFEU[108] or of whether an agreement reveals a restriction of competition by object under Article 101 TFEU.[109]

Given the limited probative value of asserted subjective motivations and the difficulties involved in establishing actual intent, the identifications of objectives will generally have to rely primarily on objective evidence concerning how a particular recommendation design operates and what effects it produces. Antitrust practice illustrates this approach. In the EU Google Shopping proceedings, for example, Google framed its defense in subjective terms, arguing that it had introduced grouped product results to improve the quality of its search service rather than to drive traffic to Google Shopping, its

---

rather than its position in individual queries and test whether the platform's own offers obtain higher visibility than comparable third-party offers, with comparability defined by attributes that may legitimately affect ranking. For evidence on the treatment of own offers, exploiting stockouts, Chen & Tsai (2024); for evidence on the rank positions of own offers, Farronato, Fradkin, & MacKay (2023).

[108] See, paradigmatically, the adjudication on predatory pricing: ECJ, 3.7.1991, Case C-62/86, *AKZO v Commission*, ECLI:EU:C:1991:286, para. 72 (holding that prices below average total costs but above average variable costs are abusive where they are set as part of a plan to eliminate a competitor, the assessment thus turning on the dominant undertaking's costs and strategy); see also ECJ, 2.4.2009, Case C-202/07 P, *France Télécom v Commission*, ECLI:EU:C:2009:214, paras. 108–110 (conforming the AKZO criteria and referring expressly to the undertaking's "eliminatory intent").

[109] See, for example, ECJ, 19.3.2015, Case C-286/13 P, *Dole Food and Dole Fresh Fruit Europe v Commission*, ECLI:EU:2015:184, paras. 117–118 (the parties' intention, while not a necessary factor, may be taken into account in determining whether an agreement is restrictive); ECJ, 18.11.2021, Case C-306/20, *Visma Enterprise*, ECLI:EU:C:2021:935, para. 69.

own comparison shopping service.[110] Both the General Court and, subsequently, the Court of Justice responded by focusing on how the design operated and on its observable effects rather than on Google's asserted motivation. Crucially, the General Court did not reject Google's account of its motivation on the ground that its true intention had been to promote its own service. Instead, it examined how the design lever operated in practice and the objective features and effects of the differential treatment between competing comparison shopping services and Google Shopping.[111] The Court of Justice found that Google had abused its dominant position in the general search market "in order to gain competitive advantages on the downstream market for specialized product search services."[112] Despite its seemingly subjective formulation, however, the court's reasoning ultimately rested on the objective function of the design lever at issue. Given the characteristics of the market and the specific circumstances of the case, the combination of the more favorable presentation of Google's own results and the demotion of competing results was found to be discriminatory and to fall outside the scope of competition on the merits.[113] Importantly, the three "specific circumstances" on which that assessment rested do not concern Google's subjective intentions. They concern, instead, the importance of the traffic generated by Google Search for comparison shopping services, users' behavior when conducting online searches, and the lack of effective substitutes for traffic lost from Google Search.[114] Equally telling in this respect is the court's objectively framed finding that traffic from Google's general results pages was redirected in favor of Google's own comparison shopping service and to the detriment of competing comparison shopping services.[115]

At this point, one might object that a taxonomy of objectives is of limited use for drawing normative implications if the pursuit of a particular objective can itself be established only by reference to the observed effects of recommendation design. Does this not introduce a degree of circularity, as adverse effects on the welfare of a particular user group are taken as evidence that a particular objective underlies the recommendation design, and that objective in turn provides the basis for regulatory intervention aimed at protecting the same user group? If actual effects provide the evidence for identifying the objective, why not attach the normative assessment directly to those effects?

The answer lies in separating the evidentiary from the normative inquiry. The two address conceptually distinct questions. The first is evidentiary: Can a particular design lever plausibly be explained by the pursuit of a particular objective? Observed effects may provide evidence for answering that question. The second is normative: What follows from the pursuit of that objective for the assessment of the recommendation design? Because the taxonomy of objectives is descriptive, this second step remains

[110] Presented in these terms in ECJ, 11.10.2024, Case C-48/22 P, *Google and Alphabet v Commission (Google Shopping)*, ECLI:EU:C:2024:726, para. 39; see also General Court, 10.11.2021, Case T-612/17, *Google and Alphabet v European Commission*, ECLI:EU:T:2021:763, para. 122.

[111] See, in particular, General Court, 10.11.2021, Case T-612/17, *Google and Alphabet v European Commission*, ECLI:EU:T:2021:763, paras. 369, 566–578.

[112] ECJ, 11.10.2024, Case C-48/22 P, *Google and Alphabet v Commission (Google Shopping)*, ECLI:EU:C:2024:726, para. 185.

[113] ECJ, 11.10.2024, Case C-48/22 P, *Google and Alphabet v Commission (Google Shopping)*, ECLI:EU:C:2024:726, para. 187.

[114] ECJ, 11.10.2024, Case C-48/22 P, *Google and Alphabet v Commission (Google Shopping)*, ECLI:EU:C:2024:726, paras. 140–146, where the Court of Justice recounts the General Court's reasoning, which it subsequently endorses and defends against the objections raised (paras. 169–172).

[115] ECJ, 11.10.2024, Case C-48/22 P, *Google and Alphabet v Commission (Google Shopping)*, ECLI:EU:C:2024:726, para. 244.

analytically distinct from the first. Using observed effects as evidence that a platform pursues a particular objective therefore does not make the normative assessment of that objective circular.

There is a further reason why the distinction matters. As the following analysis will show, identifying the objective pursued may support normative conclusions about certain constellations or scenarios on the basis of the incentives and likely effects associated with that objective, without requiring those effects to be established empirically in every individual case.

3. Regulatory Intervention Under Residual Uncertainty About Platform Objectives

A rule-maker—whether a legislature or, within the applicable legal framework, a court or administrative authority—need not always conclusively identify the objective pursued in the individual case before regulatory intervention is warranted. It may be sufficient that empirical evidence, such as an observed design lever and its effects, when assessed in light of the taxonomy and the plausible relationships it identifies between designs, objectives, and effects, gives rise to concerns with respect to particular regulatory goals. From an error-cost perspective, certain risks may reasonably be considered sufficiently serious to justify intervention even where it cannot be established with certainty that the platform in fact pursued a particular objective in the individual case.

This logic underlies, at least in part, the regulatory approach taken the risks associated with gatekeepers leveraging their control over ranking to favor their own services, including through platform-envelopment strategies, but also as facilitating straightforward and resource-efficient enforcement by avoiding the need to measure and balance welfare effects across different groups of users.[116]

This, in turn, points to a further contribution of the taxonomy: Its analytical framework can help decipher the rationale underlying particular regulatory interventions. By making explicit the objectives and risks to which certain forms of recommendation design may give rise, the taxonomy may also assist legal decision-makers, especially courts and authorities, in interpreting and applying the relevant rules. In this spirit, the next section shows how the DMA addresses recommendation design through targeted ex ante obligations responding to certain practices and the objectives and risks associated with them, rather than through a general prohibition of biased recommendations.

IV. RECOMMENDATION POWER UNDER THE DMA

The DMA is not specifically designed as an instrument for regulating recommendation power. Yet the concept provides a useful analytical lens through which to understand important elements of its regulatory logic. This is apparent from the conditions under which the DMA intervenes. It constrains platform conduct where three cumulative conditions are met: First, the conduct must be attributable to a designated gatekeeper; second, it must relate to a designated CPS; and, third, it must fall within the scope of one of the obligations exhaustively set out in Articles 5, 6, or 7 DMA. Article 13 DMA complements these obligations by preventing their circumvention.[117] Recall that we defined recommendation power as arising, in essence, where intermediation power and information power coincide. Transposed to the DMA, the point can be stated simply: Its designation regime identifies platforms exercising intermediation power, while its substantive obligations constrain, among other

[116] See Franck & Hartl (2025), 486–487.
[117] For a framework on the scope of application of Article 13 DMA, see Franck and Peitz (2024).

practices, the exercise of information power through recommendation design. Where the two coincide, the DMA effectively regulates what we conceptualize as recommendation power.

This is of considerable practical importance because recommendation design is central to many platforms covered by the DMA. The DMA's reach, however, extends beyond recommender platforms that themselves fall within its scope. Where such platforms operate as business users of gatekeepers' CPSs, they may themselves be significantly affected by the DMA and may be among the parties whose access, visibility, or commercial conditions the relevant obligations seek to protect.

After a brief discussion of the DMA's "recommendation vocabulary" in Section IV.A, Section IV.B substantiates the first part of this claim by showing how the DMA relies on intermediation power in defining its scope of application. It further demonstrates that elements of recommendation power may already play a material role at the designation stage: indeed, not only in the qualification and delineation of a CPS but, in some cases, even in determining whether the CPS constitutes an important gateway. Section IV.C then examines whether, and to what extent, particular DMA obligations constrain recommendation design.

A. On the DMA's Recommendation Vocabulary

Although various DMA obligations are clearly formulated with recommendation design in mind, the DMA contains no definitions of "recommendation" or "recommender system." This contrasts with the Digital Services Act (DSA),[118] which expressly defines a recommender system as "a fully or partially automated system used by an online platform to suggest in its online interface specific information to recipients of the service or prioritise that information, including as a result of a search initiated by the recipient of the service or otherwise determining the relative order or prominence of information displayed."[119] The DMA itself refers to "recommendations" only in defining online social networking services, which enable users to "discover other users and content."[120] The term therefore does not perform any general regulatory function in the DMA.

The operative concept most relevant for the regulation of recommendation design in the DMA is instead "ranking." Article 2(22) DMA defines ranking broadly by reference to the "relative prominence" given to goods or services or the "relevance" given to search results, irrespective of the technological means used and even where only one result is presented or communicated. The definition is broad in terms of the forms of prominence it covers but limited in terms of the CPSs to which it applies. It expressly covers online intermediation services,[121] online search engines, online social networking services, video-sharing platform services, and virtual assistants; indeed, these are arguably the types of CPS for which recommendation design is particularly central. This limitation is noteworthy because recommendation design may also matter for other CPSs, including operating systems, web browsers, and online advertising services, even though it is less central to their core functionality.

[118] Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act). OJ L 277, 27.10.2022, pp. 1–102.
[119] Article 3(s) DSA.
[120] Article 2(7) DMA.
[121] The term includes notably online marketplaces, software application stores, and accommodation and travel platforms.

Recital 52 furthermore underscores the functional breadth of the DMA's concept of "ranking." In the context of the prohibition of self-preferencing, it explains that ranking covers "all forms of relative prominence," including display, rating, linking, and voice results, as well as cases in which only a single result is presented or communicated to the end user. Because ranking encompasses voice results and single-result interfaces, even highly concentrated exercises of recommendation power, such as a single answer provided by a virtual assistant or the selection of a single prominently displayed offer, as in a "buy box," can fall within the concept. In the terminology developed in Section II, the DMA's definition tracks the economics: What is regulated is the allocation of prominence, not a particular list format.

B. Recommendation Power in the DMA's Designation Framework

Recommendation power is not, as such, a criterion for designation under the DMA, nor does the DMA require a CPS to incorporate a recommender system. The economic characteristics by which the DMA identifies gatekeepers and their CPSs, however, overlap substantially with those underlying the intermediation-power component of recommendation power (subsection 1 below). The information-power component becomes most visible at the level of the substantive obligations. Yet it also plays a role at the designation stage: in the qualification and delineation of CPSs and, as individual designation decisions illustrate, even in determining whether a particular CPS in fact functions as an important gateway (subsection 2).

1. The Intermediation-Power Component

Under Article 3(1)(b) DMA, an undertaking can be designated as a gatekeeper only if it provides a CPS "which is an important gateway for business users to reach end users." This requirement, also emphasized in Recital 15 DMA, closely corresponds to what we have conceptualized as intermediation power: Business users depend on an intermediary to reach a sufficiently large number of end users. This understanding of the gateway criterion as reflecting intermediation power is also borne out by the rationale underlying the quantitative presumption in Article 3(2)(b) DMA. Recital 20 DMA explains that, where a very large number of business users depend on a CPS to reach a very large number of end users, the provider is placed in a position "to influence the operations of a substantial part of business users to its advantage." The General Court expressly endorsed this understanding of the gateway criterion in the *ByteDance* case. It held that dependence does not require the CPS to be the only channel through which business users can reach end users; it is sufficient that the CPS constitutes an important channel for doing so. In this sense, the gateway criterion captures control over an important route of access to end users rather than the absence of alternative routes altogether.[122] Conversely, *Opera Norway* confirms that meeting the quantitative thresholds does not suffice where the CPS, in the circumstances in which it operates, lacks sufficient practical importance as a route through which business users reach end users.[123]

---

[122] General Court, 17.7.2024, Case T-1077/23, *ByteDance v Commission*, EU:T:2024:478, paras. 209–212, in particular para. 210.

[123] General Court, 2.9.2026, Case T-357/24, *Opera Norway v Commission*, ECLI:EU:T:2026:520, paras. 71–72, 84–87; see also paras. 18–22. Accordingly, for designation in respect of Microsoft Edge, the presence of elements of "information power" such as preinstallation, default settings, and other design levers that Microsoft could use to steer Windows users towards Edge, was not sufficient to establish that Edge constituted an important gateway within the meaning of Article 3(1)(b) DMA.

The characteristics underlying intermediation power also overlap substantially with those that informed the legislature's selection of the categories of services capable of constituting CPSs. Recital 13 DMA singles out "widespread and commonly used digital services that mostly directly intermediate between business users and end users" and identifies, among their characteristic features, strong network effects, multisidedness, lock-in effects, a lack of multihoming, and vertical integration.

The same logic is visible in the qualitative criteria governing designation where the quantitative thresholds do not establish the statutory presumption. Under Article 3(8) DMA, the Commission must take into account, inter alia: the number of business users using the CPS to reach end users; network effects and data-driven advantages; scale and scope effects; business-user and end-user lock-in; and vertical integration. These factors capture different mechanisms through which an intermediary may become difficult to bypass, and business users may consequently depend on it for effective access to end users.

Finally, it is worth noting that the DMA itself expressly invokes the notion of "intermediation power." Recital 26 explains that a CPS may be prone to tipping where, once its provider has obtained an advantage over rivals or potential challengers "in terms of scale or intermediation power," its position may become unassailable, making such services a particular concern for the DMA's protection of contestability. The terminology is thus close to our analytical framework. While intermediation power is not itself a separate legal criterion for designation, the DMA's designation architecture thus rests to a considerable extent on the economic conditions that generate it.

2. The Information-Power Component

The designation criteria laid down in the DMA do not suggest that information power, understood here as control over what users notice, compare, and choose, is decisive for the designation of a gatekeeper or a CPS. Yet information power does enter into the Commission's designation analysis, and the General Court's emerging case law confirms that some of the underlying mechanisms can be relevant to designation.

A particularly clear example is the Commission's analysis of TikTok's recommendation algorithm in the designation of ByteDance.[124] The Commission emphasized that TikTok's algorithm analyzed user behavior and preferences and generated personalized recommendations through which users ultimately "discover a wide range of content, creators, communities and products, thereby creating social interactions."[125] These interactions are central to the functioning of the platform: a user's "like" feeds into the algorithm and thereby affects the "recommendations and discoverability of videos."[126]

At one level, the Commission relied on these features to classify the CPS. They supported its conclusion that TikTok functions as an online social networking service rather than merely as a video-sharing platform service.[127] Information power, exercised through the recommendation algorithm, thus contributes to the legal characterization and delineation of the CPS that is subject to designation.

---

[124] European Commission, 5.9.2023, Case DMA.100040, *ByteDance – Online social networking services*, C(2023) 6102 final.
[125] European Commission (n 124), para. 45.
[126] European Commission (n 124), para. 51.
[127] European Commission (n 124), para. 42, referring to paras. 43–65.

More importantly for present purposes, however, the Commission also relied on the central role of TikTok's recommendation algorithm when rejecting ByteDance's attempt to rebut,[128] on the basis of business-user multihoming, the presumption that TikTok constitutes an important gateway within the meaning of Article 3(1)(b) DMA. The Commission's reasoning is, in essence, that business users seeking to achieve reach on TikTok must adapt their content to the platform's specific recommendation algorithm, which governs the visibility and discoverability of content. This creates platform-specific dependence and switching costs even where the same business users are also active on competing online social networking services.[129] ByteDance's information power over recommendation and discoverability thus becomes part of the explanation for why intermediation power persists despite multihoming. In terms of our analytical framework, this is precisely where the two components come together: Control over the allocation of visibility reinforces business users' dependence on the platform as a gateway to end users. The Commission's designation analysis thus provides a concrete illustration of recommendation power as defined in Section II: The platform's control over visibility becomes a mechanism through which it governs the quality of business users' effective access to end users.

The General Court subsequently confirmed this part of the Commission's reasoning. In particular, it accepted that content creators may face switching costs when shifting their activities to another platform because they must adjust their content to the different platforms' formats and algorithms, thereby confirming the relevance of platform-specific algorithmic dependence to the gateway analysis.[130]

The Commission's reasoning also reveals the rule-setting dimension of recommendation power. By recognizing that business content providers must adapt their conduct to the platform's recommendation algorithm to obtain reach, the Commission implicitly identified the recommendation system as a mechanism that structures—and, in this sense, sets the rule of—competition among business users for end users' attention. ByteDance need not prescribe to content creators what they should produce. By allocating visibility and discoverability, the algorithm instead determines the payoff structure to which creators must respond. Recommendation design thereby shapes not only what users see but also the competitive environment in which business users seek to reach them.

The Booking.com designation decision provides a related illustration that concerns the delineation of the relevant CPS.[131] The Commission considered Booking.com to constitute a single online intermediation CPS across different categories of travel services, emphasizing that business users use the service to reach the "widest possible end user base" interested in travel options.[132] This raised the question of whether the delineated CPS should also encompass the "Flight + Hotel" offering accessible through Booking.com, which linked users to a package-travel service hosted and operated by the independent third party lastminute.com. The Commission excluded that offering from the Booking.com CPS. In doing so, it stressed, among other factors, that lastminute.com, rather than Booking.com, controlled the "display and ranking of search results," while Booking.com neither had access to the

[128] See Article 3(5) DMA.
[129] European Commission (n 124), paras. 137–138.
[130] General Court, 17.7.2024, Case T-1077/23, *ByteDance v Commission*, EU:T:2024:478, paras. 163–215, in particular paras. 179 and 187; appeal pending in Case C-627/24 P, *ByteDance v Commission*.
[131] European Commission, 13.5.2024, Case DMA.100019, *Booking – Online intermediation services*, C(2024) 3176 final.
[132] European Commission (n 131), paras. 48–51.

relevant business- and end-user data nor controlled the service's front end or back end.[133] The decision thus illustrates how information power can help to identify the boundaries of intermediation power: The mere accessibility of an offering into the user journey does not make it part of the gatekeeper's CPS where control over the allocation of visibility rests with an independent provider.

C. DMA Obligations Constraining Recommendation Design

The designation framework examined above limits the DMA's intervention to gatekeepers whose CPSs exhibit the characteristics of intermediation power. The following analysis examines how particular DMA obligations constrain recommendation design and thus the exercise of information power by designated gatekeepers. In doing so, these obligations regulate the conjunction of intermediation and information power that we conceptualize as recommendation power.

1. Article 6(5) DMA: The Self-Preferencing Prohibition

Article 6(5) DMA provides that gatekeepers must not treat their own products or services more favorably in ranking, and related indexing and crawling, than similar products or services of third parties also operating on the CPS, and that gatekeepers must apply transparent, fair, and nondiscriminatory (TFN) conditions to that ranking. The rule addresses the gatekeeper's dual role as intermediary and competitor: Control over ranking and prominence enables a vertically integrated gatekeeper to favor its own offerings and thereby undermine the contestability of competing products or services.

In terms of the taxonomy of platform objectives, Article 6(5) DMA is most directly relevant to own-offer monetization (category 3), and further to strategic self-preferencing serving ecosystem dependence (category 11), dynamic foreclosure, and gateway preservation (category 12). Article 6(5) thus imposes an ex ante prohibition on one particular exercise of recommendation power: the use of control over ranking and prominence to favor the gatekeeper's own offerings.[134]

By contrast, Article 6(5) DMA does not address biases in recommendation design that favor some third-party offerings over others. It therefore does not constrain the pursuit of objectives such as the monetization of third-party transactions (category 2) or sponsored recommendations (category 5), insofar as these involve preferential treatment among third-party offerings rather than self-preferencing.

This limitation is explicit in the first sentence of Article 6(5) DMA but it also extends to its second sentence, which requires gatekeepers to apply "transparent, fair and non-discriminatory conditions to such ranking." Both the reference to "such ranking" and the explanation in Recital 52, 2nd sentence, DMA indicate that this requirement remains tied to the self-preferencing context addressed by the first sentence. Reading the second sentence as an autonomous prohibition of discrimination among third-party offerings would effectively introduce, through the back door, a substantially broader regulation of recommendation design.[135] Such an interpretation is difficult to reconcile with the DMA's regulatory approach of imposing specifically defined and legislatively endorsed "dos and don'ts."[136] A general

[133] European Commission (n 131), fn. 81 to para. 49.
[134] See Recitals 51–52 DMA; Peitz (2022, 2023); Franck & Peitz (2024), section 5.2.
[135] For an argument pointing toward a broader reading of the second sentence of Article 6(5) DMA, see Höppner (2024), paras. 410–420.
[136] To similar effect, Göhsl & Zimmer (2025), para. 142 (arguing that the provision appears deliberately to have been confined to self-preferencing).

prohibition of discrimination among third-party offerings would constitute a significantly more intrusive regulatory intervention than a prohibition of self-preferencing and would require a distinct normative and economic justification.

Although the Commission has occasionally described the reach of Article 6(5) DMA in broader terms,[137] its *Google Search* decision expressly supports this narrower understanding. The Commission states that the obligation "addresses favouring of Alphabet's distinct or additional services in ranking over similar services of third parties, and not favouring of any group of business users of the OSE over another." Significantly, the Commission adds that the interests of different groups of business users must instead be taken into account under Article 6(12) DMA, which requires fair, reasonable and non-discriminatory conditions of access to designated online search engines.[138] The Bundeskartellamt likewise proceeded on the basis of this narrow reading of Article 6(5) DMA in its Amazon decision concerning "price control mechanisms."[139]

The limits of Article 6(5) also extend to the allocation of prominence between sponsored and organic results. The mere fact that advertising modules occupy positions that are more prominent than organic search results does not by itself constitute prohibited self-preferencing. The DMA itself proceeds on the premise that CPSs may be financed by advertising; Article 6(5) DMA becomes relevant only where sponsored placement is itself used to favor the gatekeeper's own products or services.[140]

The economic rationale for prohibiting self-preferencing lies primarily in the protection of contestability in markets adjacent to the CPS. A vertically integrated gatekeeper can use its control over discovery and visibility to steer end users toward its own offerings, including where, at given prices, end users would prefer a third-party offering.[141] A familiar Chicago-style objection naturally follows: Why would a profit-maximizing platform distort recommendations rather than simply charge third-party sellers a higher fee while leaving recommendations unbiased? This objection rests, however, on the assumption that fees provide a fully flexible instrument for monetizing third-party offerings. Where the platform cannot freely adjust those fees, the incentive to self-preference may be particularly strong. This is most obvious where the fee is zero by construction, as with organic search results in *Google Shopping*, but it applies more generally where fees are uniform across broad categories of products or services or constant over time and therefore cannot reflect the platform's product-specific gains from steering users toward its own offerings. Where the gatekeeper cannot fully appropriate through fees some of the extra value generated by third-party offerings (beyond the value of the platform's own offering), preferential

[137] See European Parliament, Parliamentary Question E-000702/2024(ASW), Answer given by Mr Breton on behalf of the European Commission, 13.5.2024 ("[Article 6(5) of the [DMA] bans gatekeepers from treating its own services and products more favourably, in ranking … than similar services or products of third parties. Article 6(5) of the DMA *also* requires that ranking shall be transparent, fair and non-discriminatory" (emphasis added). The word "also" could suggest a certain degree of autonomy for the second sentence. On the other hand, the Commission expressly frames the provision as a whole in terms of "ending unfair self-preferencing practices."
[138] European Commission, 23.7.2026, Case DMA.100193 – *Alphabet – Online Search Engine – Google Search – Art. 6(5)*, C(2026) 5358 final, para. 89.
[139] Bundeskartellamt, 4.2.2026, Case No. B2-73/20, *Amazon Marketplace—Price-Control Mechanisms*, paras. 884–885. Amazon has challenged the decision before the German Federal Court of Justice (Bundesgerichtshof); proceedings are pending.
[140] See Franck & Hartl (2025), 480.
[141] Peitz (2022); Feasey & Krämer (2019); on empirical approaches, Farronato, Fradkin, & MacKay (2023); Reimers & Waldfogel (2023).

visibility and thus biased recommendation provide an alternative mechanism of value capture for the platform.[142]

As the taxonomy has shown,[143] recommendation design that appears to constitute self-preferencing aimed at own-offer monetization may in fact form part of an equilibrium-oriented strategy that ultimately benefits end users. Establishing this in an individual case, however, is likely to be difficult. Article 6(5) DMA does not provide for such a qualification of the self-preferencing prohibition. This may, in individual cases, entail a loss of consumer welfare, but it is consistent with the DMA's broader regulatory approach. Relatively clearly defined "dos and don'ts," coupled with narrowly circumscribed possibilities for justification, are intended to keep implementation and enforcement costs low; individualized case-by-case welfare analysis is deliberately avoided. Moreover, because the DMA targets only the CPSs of designated gatekeepers, its approach reflects a precautionary stance toward the risks associated with the exercise of gatekeeper power. This can be understood in error-cost terms: Where self-preferencing creates a risk of an envelopment strategy, the regulatory framework prioritizes avoiding false negatives, accepting in return a greater risk of false positives and the consumer-welfare losses they may occasionally entail.[144]

2. *Beyond Self-Preferencing: Other DMA Constraints on Recommendation Design*

Beyond Article 6(5) DMA, which directly constrains recommendation design through its rules on ranking, the DMA constrains recommendation power through two further channels: data and transparency duties, and rules constraining the gatekeeper's control over user choice, defaults, and access. Table 1 summarizes the relevant provisions and shows how each bears on recommendation power.

Table 1: DMA Provisions Relevant to Recommendation Power

Relevant DMA Provision	What It Regulates	Relevance to Recommendation Power
Art. 5(2)	Cross-service combination and use of personal data, including for advertising	Data inputs for personalization, profiling, and targeting
Arts. 5(3)–(5)	Parity clauses; restrictions on business users' steering and off-platform contracting	Recommendation design may reinforce restrictions on off-platform steering, including through reduced visibility or demotion)
Arts. 5(7)–(8)	Tying of payment, identification, browser, and other CPSs	Recommendation design may steer users toward the gatekeeper's own service stack
Arts. 5(9)–(10), 6(8)	Advertising transparency and measurement	Transparency concerning sponsored recommendations and their performance
Arts. 6(3)–(4)	Defaults, choice screens, uninstallation, and sideloading	Default and choice architecture as mechanisms for steering user choice
Arts. 6(9)–(11)	Portability and access to business-user and search data	Data inputs for competing recommendation and discovery services

[142] Franck & Peitz (2024), section 5.2.1, drawing on Peitz (2022); see also Hagiu, Teh, & Wright (2022); de Cornière & Taylor (2014). As discussed under category 3, an economic assessment is complicated by the fact that other mechanisms may be at play that may have different implications.
[143] See above III.B.Category 3: Own-Offer Monetization.
[144] See Franck & Hartl (2025), 486–487.

Arts. 6(12)	FRAND conditions of access to app stores, search engines, and social networking services	Ranking and visibility as potential elements of FRAND access conditions
Art. 13(4), (6)	Anticircumvention; nonneutral exercise of choice architecture	Ranking, friction, defaults, warnings, and interface design as potential means of circumvention

Three observations help to organize the provisions relevant for recommendation design. The first is the *data channel*, which that operates in two directions. On the one hand, Article 5(2) DMA constrains the cross-service combination of personal data on which fine-grained personalization and targeting may rely, while Article 15 DMA requires an audited description of consumer-profiling techniques. These provisions constrain or expose the data inputs underlying the gatekeeper's own recommendation rather than regulating recommendation quality as such. The Commission's Meta decision under Article 5(2) DMA illustrates the first mechanism, as the Commission emphasized that combining large volumes of personal data across services enables personalized advertising.[145]

On the other hand, Articles 6(9)–(11) DMA address access to data that may serve as inputs into rivals' competing or complementary recommendation and discovery services. Recommendation quality may depend on query, click, view, ranking, and transaction data, and Article 6(11) DMA addresses this relationship most directly by giving third-party search engines fair, reasonable, and nondiscriminatory (FRAND) access to specified data generated on the gatekeeper's search engine. The Commission's recent specification decision concerning Google Search makes the underlying rationale particularly explicit: The specified ranking data includes information on the position, prominence, and visibility of search results, while the Commission expressly linked Google's extensive access to user data to its ability to improve its search algorithm.[146] In this respect, the DMA addresses recommendation power indirectly by improving rivals' ability to effective discovery of alternatives.[147]

The second observation is the *choice-architecture channel*: Defaults, prompts, warnings, and choice screens are not always recommendations in a narrow sense, but they can perform the same steering function by making some products, services, channels, or data paths more salient or easier to choose than others. Articles 6(3) and (4) (defaults, choice screens, uninstallation, and sideloading), 6(6) (switching), and 13(6) (nonneutral choice architecture) constrain the gatekeeper's ability to use such design features to frustrate the effective exercise of choices protected by the DMA. This is illustrated by the Commission's proceedings concerning Apple's implementation of Article 6(3) DMA, in which it also examined the design of Apple's browser choice screen, default settings, and uninstallability in light of Article 13(4) and (6) DMA. Following changes including easier browser selection, adjustments to the positioning of browser icons, and expanded control over default settings, the Commission closed the proceedings.[148]

The third is the *access channel*: Article 6(12) DMA requires FRAND conditions of access for business users to app stores, search engines, and social networking services. While the provision does not itself

[145] European Commission, 23.4.2025, Case DMA.100055 – *Meta – Article 5(2)*, C(2025) 2091 final, para. 30; 72–73.
[146] European Commission, 16.7.2026, Case DMA.100209 – *SP – Alphabet – Article 6(11)*, C(2026) 5091 final. Summary of the Decision, sections 1 and 4(2), in particular point (c) on ranking data.
[147] Franck & Peitz (2024), section 3.2.
[148] European Commission, 23.4.2025, Case DMA.100185 – *Apple – iOS – Article 6(3)*, C(2025) 2562 final, paras. 5–9.

establish a right of access,[149] its FRAND requirements also apply where the relevant conditions of access include rules governing ranking. To that extent, Article 6(12) DMA operates alongside Article 6(5) DMA[150] and, unlike the latter, constrains the gatekeeper's ability to differentiate among business users. It does not, however, require identical treatment: Differentiation remains permissible where it is based on objective criteria and is objectively justified. Such differentiation may be warranted, in particular, to protect users and the platform against illegal, unsafe, deceptive, or otherwise harmful content, thus pursuing the platform-integrity objective identified as category 16 of our taxonomy.

The DMA does not, however, recognize a general integrity or safety defense applicable across its obligations. Rather, several provisions expressly provide for narrowly defined safeguards, and the Commission has treated their presence as relevant to the interpretation of obligations that contain no equivalent exception. In its Apple steering decision, for example, the Commission inferred from the absence of a security exception in Article 5(4) DMA, in contrast to Articles 6(3), 6(4), and 6(7) DMA, that the legislature had not considered implementation of Article 5(4) DMA to create security risks warranting such an exception. The Commission nevertheless left open whether objectively necessary and proportionate restrictions might still be justified by genuine end-user security concerns.[151]

There is good reason for such caution, as our taxonomy highlights the underlying identification problem: The same design lever may serve a genuine integrity objective or strategically protect the gatekeeper's own position. Measures such as warning screens, degraded interoperability, or ostensibly quality-based differentiation may therefore themselves undermine rights protected by the DMA. The relevant discipline is therefore not the mere invocation of safety but the requirements the DMA attaches to such measures, which may include necessity, proportionality, equal application, and neutral user choices. Article 13(4) and (6) DMA further prevents safety rationales from being implemented in ways that circumvent the substantive obligations.[152]

3. *Mapping the Taxonomy onto the DMA*

Bringing Sections III and IV together, Table 2 maps the taxonomy onto the DMA. The mapping identifies where the objectives pursued through recommendation design intersect with DMA obligations; it is analytical rather than a "checklist." For purposes of the mapping, we distinguish four types of DMA relevance: *direct*, where the DMA regulates conduct through which the objective is pursued; *indirect*, where it affects relevant inputs, conditions, or justifications; *conditional*, where only particular manifestations of the objective fall within DMA obligations; and *limited*, where the intersection with the DMA is narrow or peripheral.

Table 2: DMA Relevance by Taxonomy Category

Objective Category	Type of DMA Relevance	Main DMA Provisions Constraining or Regulating the Pursuit of the Objective

[149] See Recital 62, last sentence DMA.

[150] Heinz (2024), para. 111.

[151] European Commission, 23.4.2025, Case DMA.100109 – *Apple – Online Intermediation Services – app stores – AppStore – Art. 5(4)*, para. 63. An action for annulment against the decision is pending before the General Court, Case T-438/25, *Apple v Commission*.

[152] See Franck & Peitz (2024), section 5.1.2 (discussing Apple's warning screens in the context of the Dutch competition authority's decision to order Apple to adjust its App Store conditions to enable dating-app providers to use in-app payment systems other than Apple's In-App Purchase).

(2) Third-party transaction monetization	Conditional	Art. 6(12) where ranking forms part of access conditions; Art. 13 where third-party steering circumvents another DMA obligation
(3) Own-offer monetization	Direct	Art. 6(5); Art. 13
(4) End-user surplus extraction	Indirect	Art. 5(2) limits certain data inputs for personalized differentiation; Art. 6(5) where pursued through self-preferencing
(5) Sponsored recommendation	Conditional	Art. 5(9)–(10); Art. 6(8) (advertising transparency); Art. 6(5) where sponsored placement entails self-preferencing; Art. 13
(6)/(7) Attention and data monetization	Direct	Art. 5(2); Art. 6(9)–(11); Art. 6(8); Art. 15
(8) Engagement and retention	Limited	Arts. 5(2), 15 (data inputs); Arts. 6(6), 6(9), 6(13) (switching and exit); Art. 13(6)
(9) Perception shaping	Limited	Not covered as such; Art. 6(5) where perception shaping is implemented through forms of prominence constituting ranking within Art. 2(22) and involves self-preferencing
(10) Operational demand steering	Indirect	No specific obligation; operational considerations may inform whether differential treatment is objectively justified; Art. 13 where operational design circumvents a DMA obligation
(11) Ecosystem dependence	Direct	Art. 5(7)–(8), read with Art. 13; Art. 6(3)–(7)
(12) Preserving the gateway	Direct	Art. 5(3)–(5); Art. 6(5); Art. 6(11)–(13); Art. 13
(13)–(15) Persuasion, influence, and public interest	Limited	No specific DMA provisions; primarily addressed under the DSA (Arts. 27, 34–35, 38)
(16) Platform integrity and safety	Indirect	Express integrity and safety safeguards in particular obligations (e.g., Arts. 6(4), 6(7), 7(9)); Art. 13(4), (6)

The mapping supports three conclusions. First, the DMA's center of gravity lies in commercial and strategic objectives pursued through recommendation power: Own-offer monetization, attention and data, ecosystem dependence, and gateway preservation are the categories most directly addressed by the DMA.

Second, the mapping reveals a division of regulatory labor between the DMA and the DSA. Persuasion and public-interest objectives (categories 13–15) are largely left to the DSA, whose rules on recommender-system transparency, systemic-risk assessment and mitigation, and recommender systems of very large platforms address these concerns directly. A similar division emerges for engagement and retention (category 8): The Commission's preliminary finding that the addictive design of Instagram and Facebook breaches the DSA illustrates that engagement-driven recommender design is primarily a concern of the DSA rather than the DMA.[153]

Third, the operational and integrity objectives (categories 10 and 16) occupy a different position. The DMA does not regulate these objectives as such. Rather, they may become relevant in assessing whether differential treatment is objectively grounded or falls within an express safeguard recognized by a

[153] European Commission, 10.7.2026, Press Release IP/26/1579 "Commission preliminary finds the addictive design of Instagram and Facebook in breach of the Digital Services Act"; proceedings ongoing, findings preliminary.

particular objective. At the same time, Article 13 limits the gatekeeper's ability to invoke such objectives through the design choice that in substance circumvents a DMA obligation.

4. Article 13 DMA: Recommender Design as Circumvention

Article 13 DMA prohibits conduct that undermines effective compliance with Articles 5, 6, and 7 DMA. It is particularly relevant to recommendation design because a gatekeeper may respond to a substantive prohibition by shifting ranking, defaults, friction, warnings, or other design features that preserve effects of the prohibited conduct in a different form. The circumventing practice need not therefore take the same form as the conduct addressed by the substantive obligation; what matters is whether it undermines effective compliance with that obligation. Article 13(4) DMA explicitly extends to behavior of a contractual, commercial, technical, or any other nature, including behavioral techniques and interface design; Article 13(6) DMA captures, inter alia, degrading the conditions or quality of services for those who exercise DMA rights, presenting of choices in a nonneutral manner, and using interface design to subvert end-user or business-user autonomy.[154]

In earlier work, we proposed a three-step, effects-based methodology for identifying such circumvention. First, one identifies the target effects of the relevant DMA obligation: Which contestability or fairness problem is the rule intended to address, what change in conduct does it seek to induce, and what are its expected effects? Second, one examines the market effects of the suspicious practice: How does it alter the incentives and choices of end users and business users? Third, one compares the effects of the suspicious practice with those the obligation seeks to prevent and assesses their equivalence: Are they sufficiently similar for the practice to undermine the contestability or fairness that the obligation is meant to secure? This assessment is necessarily context-specific, since equivalence may depend on the market environment and, potentially, on the position of particular business users.[155] Recommendation design is the natural habitat of this test: Once a first-order practice is prohibited, ranking, friction, defaults, warnings, and other interface features provide alternative means through which a gatekeeper may seek to preserve its effects. Three applications illustrate this point.

a) De-Ranking After the Ban on Parity Clauses

Article 5(3) DMA protects business users' freedom to offer different prices or conditions through other channels. Once parity clauses are prohibited, a hotel or seller may therefore offer a lower price outside of the platform. Yet the gatekeeper may be able to recreate similar discipline through recommendation design. If its ranking algorithm uses on-platform conversion rates, an offer may be demoted because lower off-platform prices divert transactions away from the platform and thereby reduce its on-platform conversion rate. De-ranking then raises the cost of offering better terms elsewhere and may induce sellers to maintain de facto parity. The initial problem is contractual price parity; recommendation power becomes the enforcement technology.[156]

The three-step circumvention test makes the problem transparent. Article 5(3) DMA seeks to preserve business users' ability to differentiate their terms across channels without being disciplined into parity.

---

[154] See Recital 70 to the DMA; Franck & Peitz (2024), sections 2–3.
[155] Franck & Peitz (2024), section 3.
[156] Franck & Peitz (2024), section 4.3.2; Hunold, Kesler, & Laitenberger (2020); for the underlying economics of parity clauses, Edelman & Wright (2015); Boik & Corts (2016); Johnson (2017); Wang & Wright (2020).

The suspicious practice changes form but may recreate the effect the obligation seeks to prevent: De-ranking does not formally prohibit lower off-platform prices but it might degrade the quality of platform access to the point where such differentiation becomes commercially unattractive. If so, the ranking mechanism restores through visibility what the gatekeeper can no longer impose contractually.

Article 13 DMA is therefore the most natural legal basis for addressing such conduct: It protects the effective exercise of the freedom secured by Article 5(3) DMA, rather than merely its formal contractual availability. On our earlier analysis, such de-ranking is better addressed through Article 13(4) and (6) DMA than by extending Article 5(3) DMA itself beyond its proper scope.[157]

This does not mean, however, that de-ranking based on conversion rates is inherently suspect. The distinction matters because it illustrates why an effects-based approach must examine the function of the ranking rule rather than condemn a particular algorithmic input. As discussed under category 2, a low conversion rate can reflect two different mechanisms. It may indicate poor product fit or low quality, in which case lower ranking may improve recommendations for users. But it may also result from users discovering a product on the platform and purchasing it more cheaply off the platform, in which case treating the lower conversion rate as a negative ranking signal may penalize precisely the channel differentiation protected by Article 5(3) DMA. Nor does it follow that the solution is simply to prohibit using off-platform price information. A consumer-oriented ranking may need to take such prices into account; indeed, correcting conversion rates for lower off-platform prices may be necessary to avoid systematically demoting offers that consumers value but purchase elsewhere. The relevant question is therefore not whether a particular input enters the algorithm but whether the ranking rule neutralizes the exercise of DMA-protected channel choice. Accordingly, once it is established that the CPS uses conversion rates as a ranking factor, there is a strong case for finding the requisite equivalence of effects with an explicit parity policy, unless the platform can demonstrate that it adjusts the conversion-rate signal in the seller's favor where better terms are offered off-platform.[158]

b) Sponsored Modules and First-Party Boxes

Once Article 6(5) DMA bans self-preferencing within an organic ranking, a gatekeeper may shift the allocation of prominence to other elements of the interface. It may, for example, move first-party offers out of the organic ranking and into separate labeled boxes or sponsored modules. If those modules dominate the user journey, recommendation power may thus be exercised through a different design element while producing effects equivalent to self-preferencing in the organic ranking.

Early evidence on Google's implementation of Article 6(5) DMA illustrates the potential pattern. Alongside a new box increasing the visibility of competing comparison shopping services, Google introduced a product viewer that has been characterized as functioning as Google's own comparison service. The latter was found to be displayed for a larger share of searches, to appear more often per page, and to occupy more screen space. On that interpretation, prominence did not simply recede in response to Article 6(5) DMA but shifted across interface elements.[159]

[157] Franck & Peitz (2024), section 4.3.2.
[158] Franck & Peitz (2024), section 4.3.2; Hunold, Kesler, & Laitenberger (2020).
[159] Püplichhuisen & Sirries (2026), based on more than 1.2 million product-related searches on Google's German desktop version between 2022 and 2024, with visibility measured by trigger rates, rank, number of occurrences per page, and occupied screen space. The assessment of the product viewer as functioning as Google's own comparison shopping service may be contested.

The Commission's recent non-compliance decision provides support for viewing prominence across interface elements as relevant to Article 6(5) DMA. The Commission interprets ranking as encompassing not only the position of a service in a list of search results but "any other form of relative prominence". Relevant dimensions include the frequency and amount of information displayed, the use of particular display features, such as pictures or featured units, and the possibilities for users to engage with the content.[160] In its assessment of Google's comparison-shopping services, the Commission accordingly does not ask merely whether third-party services received greater prominence following Google's compliance changes. Rather, it asks whether their prominence matches that afforded to Google's own comparison-shopping services.[161] It concludes that increased prominence for third-party comparison-shopping services through rich results does not offset the greater prominence afforded to Google's own services through dedicated units.[162] This lends support to assessing prominence across the interface rather than treating individual design elements in isolation. It does not, however, establish that a shift of prominence across interface elements constitutes circumvention under Article 13 DMA.

Under the effects-based approach developed above, the relevant question is whether the redesigned interface restores through separate modules the preferential prominence that Article 6(5) DMA seeks to prevent within ranking. If it does, Article 13 DMA may bring the redesign within the reach of Article 6(5) DMA even though the preferential treatment has migrated to a formally distinct interface element.

The assessment remains delicate, however. An ad-funded redesign may reflect a genuine change in the platform's business-model and monetization strategy may fall outside Article 13 DMA. The harder it is to establish the requisite connection between the redesign and the effects targeted by a specific DMA obligation, the weaker the case for circumvention. Article 13 DMA empowers the Commission and courts to safeguard the effectiveness of existing DMA obligations; it does not provide a free-standing mandate to redesign platform business models.[163]

c) Selected Third Parties and Service Bundles

A platform may also shift from favoring its own downstream offer to favoring selected third-party sellers or app providers—for instance by making better ranking or greater visibility conditional on their use of the gatekeeper's logistics, payment, identity, or technical services. Recommendation power then becomes an instrument of service-stack steering: Rather than directing end users toward the gatekeeper's own offer, the ranking mechanism rewards third parties that integrate into its service stack and thereby creates incentives for business users to do so.

The mere preferential treatment of selected third parties does not, on the interpretation of Article 6(5) adopted above,[164] amount to prohibited self-preferencing. Different DMA provisions may nevertheless become relevant depending on how the arrangement is structured. Article 6(5) DMA may apply where the favored third-party offer is attributable to the gatekeeper; Article 5(7) or 6(12) DMA may constrain particular forms of tying or discriminatory access conditions. Most importantly for present purposes,

[160] European Commission, 23.7.2026, Case DMA.100193 – *Alphabet – Online Search Engine – Google Search – Art. 6(5)*, C(2026) 5358 final, paras. 75–77.
[161] European Commission (n 160), paras. 248–250.
[162] European Commission (n 160), paras. 251–254; see also paras 711–713.
[163] Franck & Peitz (2024), sections 3.5 and 5.2.2; see also Peitz (2022), 10, 12; Bostoen (2023), 290, 301.
[164] See above, Section IV.C.1.

Article 13 DMA may apply where recommendation design is used to reproduce through third-party incentives effects that a substantive DMA obligation seeks to prevent. The Italian FBA case illustrates the underlying mechanism in an abuse-control setting: Preferential visibility for sellers using the gatekeeper's own logistics service can turn recommendation design into an incentive for business users to adopt another service in the gatekeeper's ecosystem.[165]

d) Summary

Taken together, the three applications illustrate the future-proofing function of Article 13 DMA with respect to recommendation power. Recommendation design can reproduce the effects of conduct prohibited by a DMA obligation even where that obligation does not itself concern ranking or recommendation. The anticircumvention test is effects-based, not form-based, and its outer limits keep the Commission and the courts from becoming product designers.[166]

V. CONCLUDING REMARKS

This article has developed recommendation power as the operational concept at the intersection of information power and intermediation power: the active, user-facing steering of choice by a platform that is also an effective-access bottleneck. What the platform makes prominent becomes what is chosen; and, where access runs through the platform, control over prominence is control over access. Because ranking and recommendation are continuous, opaque design choices with dynamic effects, recommendation power strains ex post abuse control, thereby helping to explain the turn toward ex ante intervention reflected in the DMA and similar regulatory interventions.

The taxonomy of objectives proposed in Section III is the analytical core of the article. It distinguishes an independent-advice benchmark from fifteen other objectives that may systematically shape recommendation design, ranging from immediate monetization to longer-term objectives and objectives of influence, welfare, and integrity. Crucially, it classifies objectives rather than design levers: The same ranking rule, default, label, or interface feature may serve very different purposes. The taxonomy therefore provides a framework for moving from observed recommendation design to hypotheses about the objective pursued, prevailing incentives and effects that may follow, while making explicit the uncertainty inherent in that interference.

The DMA, finally, does not name recommendation power; it nevertheless reaches it functionally and selectively. Its designation regime identifies situations in which intermediation power is sufficiently pronounced to warrant ex ante regulation, while its substantive obligations constrain particular ways in which the gatekeeper may combine that position with control over visibility and user choice. The ban on self-preferencing in ranking pursuant to Article 6(5) DMA is the clearest example: It prohibits one particular exercise of recommendation power but does not establish a general duty of unbiased recommendation or, on our reading, a general prohibition of discrimination among third-party offerings. Other provisions constrain the data inputs, choice architecture, access conditions, and business-user

[165] Autorità Garante della Concorrenza e del Mercato (AGCM), 30.11.2021, Case A528 – *FBA Amazon*, Decision No. 29925 (on the pending judicial review, see above note 80), in particular paras. 253–254, 698, 737; see, for an English-language account of the decision, AGCM, 9.12.2021, Press Release, "Italian Competition Authority: Amazon fined over €1,128 billion for abusing its dominant position." See also OECD (n 80), paras. 55, 58; Colangelo (2023), 549–550; Franck & Peitz (2024), section 5.2.2.
[166] Franck & Peitz (2024), sections 3.5 and 6.

rights through which recommendation power may be exercised. Article 13 DMA completes this architecture by preventing recommendation design from being used to reproduce the effects of conduct prohibited elsewhere in the DMA. Thus, the anticircumvention provision is a powerful future-proofing device but not a free-standing mandate to regulate recommendation design: Captured circumvention must remain tied to the effects that an DMA obligation seeks to prevent.

Under the DMA, the question is whether recommendation design interferes with a choice, access right, or competitive opportunity that the Regulation protects. This also marks the limit of the DMA as a framework for recommendation power. Many forms of recommendation bias identified by the taxonomy remain outside its scope, not because they are necessarily benign but because the DMA is a regime of specifically defined obligations rather than a general code of recommender neutrality.

The importance of that distinction is likely to grow as recommendation interfaces become more concentrated. Virtual assistants[167] and single-answer interfaces return one option rather than a ranked list, collapsing the allocation of prominence and the grant of effective access into a single act. The DMA's definition of ranking already anticipates this possibility by including the presentation of a single result. More generally, our analytical apparatus of benchmark, objective, design lever, and effects does not depend on the existence of a visible ranked list. The fewer alternatives the user sees, the more consequential the exercise of recommendation power becomes.

[167] Whether and to what extent generative AI assistants fall within the DMA's definitions of virtual assistant (Art. 2(12) DMA) or online search engine is an open question; the functional point is independent of it.

REFERENCES

Allcott, H. & M. Gentzkow (2017), Social Media and Fake News in the 2016 Election, 31 Journal of Economic Perspectives 211–236.

Allcott, H., L. Braghieri, S. Eichmeyer, & M. Gentzkow (2020), The Welfare Effects of Social Media, 110 American Economic Review 629–676.

Allcott, H., M. Gentzkow, & L. Song (2022), Digital Addiction, 112 American Economic Review 2424–2463.

Anderson, S. & S. Coate (2005), Market Provision of Broadcasting: A Welfare Analysis, 72 Review of Economic Studies 947–972.

Anderson, S. & A. de Palma (2012), Competition for Attention in the Information (Overload) Age, 43 RAND Journal of Economics 1–25.

Anderson, S. & M. Peitz (2023), Ad Clutter, Time Use, and Media Diversity, 15 American Economic Journal: Microeconomics 227–270.

Anderson, S. & M. Peitz (2026), The Two Sides of Platform Ad Load, unpublished manuscript.

Andreoli-Versbach, P. & J. Gans (2024), Interplay between Amazon Store and Logistics, 20 European Competition Journal 460–497.

Armstrong, M. & J. Zhou (2011), Paying for Prominence, 121 Economic Journal F368–F395.

Athey, S. & G. Ellison (2011), Position Auctions with Consumer Search, 126 Quarterly Journal of Economics 1213–1270.

Athey, S. & F. Scott Morton (2022), Platform Annexation, 84 Antitrust Law Journal 677–703.

Bakos, J.Y. (1997), Reducing Buyer Search Costs: Implications for Electronic Marketplaces, 43 Management Science 1676–1692.

Bar-Isaac, H. & S. Shelegia (2026), Monetizing Steering, Journal of Political Economy Microeconomics. doi: 10.1086/740100.

Belleflamme, P. & M. Peitz (2021), The Economics of Platforms: Concepts and Strategy, Cambridge University Press.

Bergemann, D. & A. Bonatti (2019), Markets for Information: An Introduction, 11 Annual Review of Economics 85–107.

Boik, A. & K. Corts (2016), The Effects of Platform Most-Favored-Nation Clauses on Competition and Entry, 59 Journal of Law and Economics 105–134.

Bordalo, P., N. Gennaioli, & A. Shleifer (2013), Salience and Consumer Choice, 121 Journal of Political Economy 803–843.

Bostoen, F. (2023), Understanding the Digital Markets Act, 68 Antitrust Bulletin 263–306.

Bourreau, M. & G. Gaudin (2022), Streaming Platform and Strategic Recommendation Bias, 31 Journal of Economics & Management Strategy 25–47.

Budaraju, H. (2025), AI Overviews Are Now Available in Over 200 Countries and Territories, and More than 40 Languages, May 20, 2025. https://blog.google/products-and-platforms/products/search/ai-overview-expansion-may-2025-update/.

Burguet, R., R. Caminal, & M. Ellman (2015), In Google We Trust? 39 International Journal of Industrial Organization 44–55.

Caldentey, R. & T. Xie (2026), Intertemporal Demand Allocation for Inventory Control in Online Marketplaces, arXiv:2604.07312.

Carlton, D. & M. Waldman (2002), The Strategic Use of Tying to Preserve and Create Market Power in Evolving Industries, 33 RAND Journal of Economics 194–220.

Carroni, E., L. Madio, & S. Shekhar (2024), Superstar Exclusivity in Two-Sided Markets, 70 Management Science 991–1011.

Che, Y.-K. & J. Hörner (2018), Recommender Systems as Mechanisms for Social Learning, 133 Quarterly Journal of Economics 871–925.

Chen, N. & H.-T. Tsai (2024), Steering via Algorithmic Recommendations, 55 RAND Journal of Economics 501–518.

Chen, Y. & C. He (2011), Paid Placement: Advertising and Search on the Internet, 121 Economic Journal F309–F328.

Choi, Y.S. & K. Fuchikawa (2024), Self-Preferencing in Korea and Japan, 47 World Competition 495–520.

Colangelo, G. (2023), Antitrust Unchained: The EU's Case Against Self-Preferencing, 72 GRUR International 538–556.

Crémer, J., Y.-A. de Montjoye, & H. Schweitzer (2019), Competition Policy for the Digital Era, Report for the European Commission.

Cui, R. & H. Shin (2018), Sharing Aggregate Inventory Information with Customers: Strategic Cross-Selling and Shortage Reduction, 64 Management Science 381–400.

Cui, R., D.J. Zhang, & A. Bassamboo (2019), Learning from Inventory Availability Information: Evidence from Field Experiments on Amazon, 65 Management Science 1216–1235.

de Cornière, A. & G. Taylor (2014), Integration and Search Engine Bias, 45 RAND Journal of Economics 576–597.

de Cornière, A. & G. Taylor (2019), A Model of Biased Intermediation, 50 RAND Journal of Economics 854–882.

Dellarocas, C. (2006), Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms, 52 Management Science 1577–1593.

Dinerstein, M., L. Einav, J. Levin, & N. Sundaresan (2018), Consumer Price Search and Platform Design in Internet Commerce, 108 American Economic Review 1820–1859.

Edelman, B. & J. Wright (2015), Price Coherence and Excessive Intermediation, 130 Quarterly Journal of Economics 1283–1328.

Eifert, M., A. Metzger, H. Schweitzer, & G. Wagner (2021), Taming the Giants: The DMA/DSA Package, 58 Common Market Law Review 987–1028.

Eisenmann, T., G. Parker, & M. Van Alstyne (2011), Platform Envelopment, 32 Strategic Management Journal 1270–1285.

Eliaz, K. & R. Spiegler (2011), A Simple Model of Search Engine Pricing, 121 Economic Journal 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, 62 Review of Industrial Organization 119–148.

Falce, V. & N.M.F. Faraone (2023), Digital Ecosystems in the Wake of a Legislative/Regulatory Turmoil: A First (Tentative) Antitrust Assessment of the Italian (and European) Experience in the AGCM Case Law, 46 World Competition 37–64.

Farronato, C., A. Fradkin, & A. MacKay (2023), Self-Preferencing at Amazon: Evidence from Search Rankings, 113 AEA Papers and Proceedings 239–243.

Feasey, R. & J. Krämer (2019), Implementing Effective Remedies for Anti-Competitive Intermediation Bias on Vertically Integrated Platforms, CERRE Report.

Fischer E., L. Hornkohl, & N. Imgarten (2025), Discriminatory Leveraging Plus: The Standard for Independent Self-Preferencing Abuses after *Google Shopping* (C-48/22 P), 10 European Papers – A Journal on Law and Integration 25–44.

Fleder, D. & K. Hosanagar (2009), Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity, 55 Management Science 697–712.

Ford, George S. (2024), Amazon: A Monopolist That Undersells Its Competitors? Phoenix Center Policy Bulletin No. 67.

Franck, J.-U. (2025), Implementing and Enforcing the EU's Digital Gatekeeper Regulation: A German Perspective on the Role of National Authorities. In: F. Decarolis et al. (eds.), The EU Digital Regulation and Its Impact on Member States, Springer Nature: Cham, 217–251.

Franck, J.-U. & M. Hartl (2025), Coping with Undesirable Effects of DMA Implementation: Google Search and the Ban on Self-Preferencing, 16 Journal of European Competition Law & Practice 472–491.

Franck, J.-U. & M. Peitz (2021), Digital Platforms and the New 19a Tool in the German Competition Act, 12 Journal of European Competition Law & Practice 513–528.

Franck, J.-U. & M. Peitz (2024), The Digital Markets Act and the Whack-a-Mole Challenge, 61 Common Market Law Review 299–344.

Fudenberg, D. & D.K. Levine (2006), A Dual-Self Model of Impulse Control, 96 American Economic Review 1449–1476.

Gauthier, G., R. Hodler, P. Widmer, & E. Zhuravskaya (2026), The Political Effects of X's Feed Algorithm, 652 Nature 416–423.

Google (2025), Search Quality Evaluator Guidelines, September 11, 2025. https://static.googleusercontent.com/media/guidelines.raterhub.com/en//searchqualityevaluatorguidelines.pdf.

Göhsl, J.-F. & D. Zimmer (2025), Artikel 6 DMA, in T. Körber, H. Schweitzer, & D. Zimmer (eds.), Immenga/Mestmäcker, Wettbewerbsrecht, Vol. 1, 7th ed., C.H. Beck: Munich.

Hagiu, A. & B. Jullien (2011), Why Do Intermediaries Divert Search? 42 RAND Journal of Economics 337–362.

Hagiu, A., T.-H. Teh, & J. Wright (2022), Should Platforms Be Allowed to Sell on Their Own Marketplaces? 53 RAND Journal of Economics 297–327.

Hagiu, A. & J. Wright (2023), Data-Enabled Learning, Network Effects, and Competitive Advantage, 54 RAND Journal of Economics 638–667.

Hagiu, A. & J. Wright (2024), Marketplace Leakage, 70 Management Science 1529–1553.

Heidhues, P., M. Köster, & B. Kőszegi (2023), Steering Fallible Consumers, 133 Economic Journal 1430–1465.

Heinz, S. (2024), Article 6(5) DMA [Prohibition of Self-Preferencing in Ranking], in R. Podszun (ed.), Digital Markets Act, Article-by-Article Commentary, Nomos: Baden-Baden.

Helberger, N., K. Karppinen, & L. D'Acunto (2018), Exposure Diversity as a Design Principle for Recommender Systems, 21 Information, Communication & Society 191–207.

Höppner, T. (2024), Self-Preferencing in Online Search under Article 6(5) DMA, Nomos: Baden-Baden.

Hosseinmardi, H., A. Ghasemian, A. Clauset, M. Mobius, D.M. Rothschild, & D.J. Watts (2021), Examining the Consumption of Radical Content on YouTube, 118 Proceedings of the National Academy of Sciences e2101967118.

Hua, X. & K.E. Spier (2025), Holding Platforms Liable, 17 American Economic Journal: Microeconomics 68–101.

Hunold, M., R. Kesler, & U. Laitenberger (2020), Rankings of Online Travel Agents, Channel Pricing, and Consumer Protection, 39 Marketing Science 92–116.

Ichihashi, S. & B.-C. Kim (2023), Addictive Platforms, 69 Management Science 1127–1145.

Janssen, M.C.W., T. Jungbauer, M. Preuss, & C. Williams (2026), Search Platforms: Big Data and Sponsored Positions, Economic Journal. doi: 10.1093/ej/ueag002.

Johnson, J. (2017), The Agency Model and MFN Clauses, 84 Review of Economic Studies 1151–1185.

Jürgensmeier, L. & B. Skiera (2025), Measuring Self-Preferencing on Digital Platforms, Journal of Marketing. doi: 10.1177/00222429251360772.

Kamenica, E. & M. Gentzkow (2011), Bayesian Persuasion, 101 American Economic Review 2590–2615.

Kleinberg, J., S. Mullainathan, & M. Raghavan (2024), The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization, 70 Management Science 6336–6355.

Lee, K.H. (2024), Abuse of Dominance in Digital Markets: The Korea Fair Trade Commission's Google Play Store Decision, 37 Seoul Journal of Economics 37–54.

Lee, K.H. & L. Musolff (2026), Two-Sided Markets Shaped by Platform-Guided Search, Econometrica, forthcoming.

Lee, S. (2025), On the Japan Fair Trade Commission's Google Decision: Some Early Reflections, Kluwer Competition Law Blog, April 29, 2025.

Luca, M. & G. Zervas (2016), Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud, 62 Management Science 3412–3427.

Mayzlin, D., Y. Dover, & J. Chevalier (2014), Promotional Reviews: An Empirical Investigation of Online Review Manipulation, 104 American Economic Review 2421–2455.

Motta, M. (2023), Self-Preferencing and Foreclosure in Digital Markets: Theories of Harm for Abuse Cases, 90 International Journal of Industrial Organization 102974.

Motta, M. & A. Penta (2026), Market Effects of Brand Search Advertising, RAND Journal of Economics, forthcoming.

Ott, M., C. Cardie, & J. Hancock (2012), Estimating the Prevalence of Deception in Online Review Communities, Proceedings of the 21st International Conference on World Wide Web (WWW 2012).

Peitz, M. (2022), The Prohibition of Self-Preferencing in the DMA, CERRE Report.

Peitz, M. (2023), How to Apply the Self-Preferencing Prohibition in the DMA, 14 Journal of European Competition Law & Practice 310–315.

Peitz, M. (2024), Digital Attention Intermediaries, 37 Seoul Journal of Economics 1–35.

Peitz, M. & H. Schweitzer (2016), Suchmaschinen unter der Lupe: Informationsherrschaft und ihre Schranken, 96 Wirtschaftsdienst 825–829.

Peitz, M. & A. Sobolev (2025), Biased Recommendations and Differentially Informed Consumers, 13(2) Economic Theory Bulletin 245–261.

Peitz, M. & Sobolev, A. (2026a), Inflated Recommendations, 57 RAND Journal of Economics 368–385.

Peitz, M. & Sobolev, A. (2026b), Recommendation Power and Competition, International Journal of Industrial Organization 103310. doi: 10.1016/j.ijindorg.2026.103310.

Persch, J. (2021), Is Taking Away User Traffic Anti-Competitive? Munich Court Stops Cooperation between Google and the German Federal Ministry of Health, 12 Journal of European Competition Law & Practice 694–697.

Püplichhuisen, P. & S. Sirries (2026), DMA in Effect and (In)Effective? Competition and Visibility in Online Product Search, Journal of Competition Law & Economics. doi: 10.1093/joclec/nhag022.

Rayo, L. & I. Segal (2010), Optimal Information Disclosure, 118 Journal of Political Economy 949–987.

Reid, E. (2025), AI in Search: Going Beyond Information to Intelligence, May 20, 2025. https://blog.google/products-and-platforms/products/search/google-search-ai-mode-update/.

Reimers, I. & J. Waldfogel (2023), A Framework for Detection, Measurement, and Welfare Analysis of Platform Bias, NBER Working Paper No. 31766.

Resnick, P. & R. Sami (2007), The Influence Limiter: Provably Manipulation-Resistant Recommender Systems, Proceedings of the 2007 ACM Conference on Recommender Systems 25–32.

Resnick, P. & R. Sami (2008), The Information Cost of Manipulation-Resistance in Recommender Systems, Proceedings of the 2008 ACM Conference on Recommender Systems 147–154.

Rhodes, A. and J. Zhou (forthcoming), Personalization and Privacy Choice, RAND Journal of Economics.

Roesler, A.-K. & B. Szentes (2017), Buyer-Optimal Learning and Monopoly Pricing, 107 American Economic Review 2072–2080.

Russell, B. (1938), Power: A New Social Analysis, Norton: New York.

Schweitzer, H., J. Haucap, W. Kerber, & R. Welker (2018), Modernisierung der Missbrauchsaufsicht für marktmächtige Unternehmen, Nomos.

Schweitzer, H. & R. Welker (2019), Competition Policy for the Digital Era, 3(2) CPI Antitrust Chronicle, 16–24.

Scott Morton, F. (2023), Improving the Contestability of E-Commerce in Two Jurisdictions: The Amazon Case, Bruegel Policy Brief, No. 22/2023, Bruegel: Brussels.

von Wangenheim, J. (2025), Optimal Information Design of Online Marketplaces with Return Rights, 133 Journal of Political Economy 3019–3043.

Vosoughi, S., D. Roy, & S. Aral (2018), The Spread of True and False News Online, 359 Science 1146–1151.

Wang, C. & J. Wright (2020), Search Platforms: Showrooming and Price Parity Clauses, 51 RAND Journal of Economics 32–58.

White, A. (2013), Search Engines: Left Side Quality versus Right Side Profits, 31 International Journal of Industrial Organization 690–701.

Xu, L., J. Chen, & A. Whinston (2012), Effects of the Presence of Organic Listing in Search Advertising, 23 Information Systems Research 1284–1302.

Zennyo, Y. (2022), Platform Encroachment and Own-Content Bias, 70 Journal of Industrial Economics 684–710.