UCLouvain

UNIVERSITÉ DE NAMUR

**Economics School of Louvain – ESL**
**Economics School of Namur - ESN**

# Self-Preferencing in Digital Markets: The Role of Behavioral Biases in the Google Search Case

Author : Sow Fatoumata
Thesis Director : Johannes Johnen
Thesis Reader : Kiedaisch Christian
Academic Year : 2025-2026
Master in Economics – 120 credits – Focus : Professional

Université catholique de Louvain - Economics School of Louvain - Place Montesquieu 3, 1348 Louvain-La-Neuve
Université de Namur – Economics School of Namur – Rempart de la Vierge 8, 5000 Namur

**Abstract**

This dissertation investigates whether Google exploits consumer behavioral biases to practice self-preferencing, or whether its interface design simply helps rational users find better results faster. Two competing hypotheses organize the analysis. Under H1, Google's choice architecture exploits limited attention, default inertia, and visual salience to generate a click advantage independent of product quality. Under H2, the same design elements reduce search costs and reflect genuine quality superiority.

The existing evidence is more consistent with H1 than H2. One third of users are persistently inattentive, active choice produces negligible switching, and visual salience captures attention before rational evaluation begins. However, no decisive experimental test exists yet. The dissertation proposes four experimental designs that would close this gap and identifies four specific limitations of the Digital Markets Act for persistently inattentive consumers, advancing two concrete policy recommendations.

**Keywords:** self-preferencing, behavioral biases, Google Shopping, Digital Markets Act, limited attention, visual salience, choice architecture, steering, search engine.

**Acknowledgements**

This thesis would not have been what it is without the guidance of Professor Johnen. His kindness, his constructive feedback, and his suggestions on the literature shaped this dissertation in ways I could not have achieved on my own. I am grateful to have had him as my thesis director.

I would also like to thank my thesis reader for taking the time to read and evaluate this work.

My deepest thanks go to my husband. He supported me throughout this entire process, believed in me when I doubted myself, and provided emotional support that was truly invaluable.

I also want to mention my son, who is two and a half years old and has no idea what a thesis is, but whose smile made every difficult moment a little easier.

My sisters deserve special recognition. They listened to me talk about this dissertation more times than anyone should have to, and they did so with patience and love. To all my brothers and sisters, thank you for being present. I am lucky to have you.

I also want to thank my friends, who made this journey so much more than just an academic experience. Their presence, their humor, and their support meant more than they probably know.

Finally, I want to express my gratitude to UCLouvain as an institution. The quality of its faculty, the richness of its academic environment, and the warmth of the people I encountered along the way have made this university a place that will stay in my heart.

AI language tools were used to assist with English writing style and clarity. All ideas, arguments, analysis, and academic judgments are the author's own.

## Table des matières


***Introduction*** *5*

***Chapter 1 — Self-Preferencing in Digital Markets: Mechanisms, Evidence, and the Limits of Structural Remedies*** *6*
1.1 The economics of gatekeeping platforms 6
1.2 The Google Shopping case: a primer 7
1.3 Forms of self-preferencing: a taxonomy 8
1.4 Why structural remedies may be insufficient 8

***Chapter 2 — From Rational Search to Behavioral Exploitation: A Review of the Literature*** *9*
2.1 Search costs, position effects, and the rational baseline 9
2.2 The personalization paradox: Kaye (2025) and the pricing response problem 10
2.3 Behavioral Industrial Organization: the theoretical core 11

***Chapter 3 — Behavioral Mechanisms: How Self-Preferencing Operates in Practice*** *12*
3.1 Limited attention 13
3.2 Default bias and the limits of choice architecture 14
3.3 Visual salience and the design of the search results page 15

***Chapter 4 — Exploitation or Information? An Analytical Framework*** *16*
4.1 The central dilemma 16
4.2 Formalising the two hypotheses 17
4.3 The observational equivalence problem 18
4.4 What the existing evidence shows 18
4.5 The central gap and proposed experimental designs 19
4.6 What this chapter establishes 21

***Chapter 5 — The Google Shopping Case Through a Behavioral Lens*** *21*
5.1 The click-through rate data as behavioral evidence 21
5.2 The rejection of Google's counterfactual studies 22
5.3 What the case establishes and what it leaves open 23

***Chapter 6 — The Digital Markets Act: Relevant Obligations and Behavioral Foundations*** *23*
6.1 Gatekeepers and the logic of ex ante regulation 24
6.2 Article 6(5) — Non-discriminatory ranking 24
6.3 Article 6(3) — Choice screens and defaults 25
6.4 Article 6(11) — Data access for competing search engines 25
6.5 The economists behind the DMA 26

***Chapter 7 — The Limits of the Digital Markets Act: What Behavioral Economics Reveals*** *27*
7.1 A necessary but insufficient framework 27
7.2 Limitation 1 — Article 6(5) regulates ranking but not visual salience 27
7.3 Limitation 2 — Choice screens do not reach persistently inattentive users 28
7.4 Limitation 3 — The DMA is designed for rational consumers 29
7.5 Limitation 4 — The effectiveness gap 30

***Conclusion*** *32*

***Bibliography*** *35*

# Self-Preferencing in Digital Markets: The Role of Behavioral Biases in the Google Search Case

## Introduction

In June 2017, the European Commission issued one of the most consequential decisions in the history of digital regulation. After seven years of investigation, it found that Google had systematically abused its dominant position in general internet search by redirecting users toward its own comparison shopping service, Google Shopping, while demoting rival services in its organic results. The fine, €2.42 billion, was unprecedented. But the legal verdict settled far less than it appeared to. The Commission documented, through click-through rate data and eye-tracking studies, that users gravitate toward the first results and that visually rich formats significantly increase engagement. It even acknowledged the role of ranking bias and saliency bias in shaping user behaviour. But it stopped short of asking the deeper question: whether these patterns reflect cognitive exploitation or simply a rational response to better-presented information.

That question sits at the heart of this dissertation.

The standard account of the Google Shopping case is structural. Google occupied a chokepoint in the information economy, exploited that position to divert traffic from competitors, and the Commission intervened to restore a level playing field. This account is legally coherent and empirically grounded. But it is, this thesis will argue, incomplete. It treats consumers as passive victims of an architectural manipulation without examining the cognitive mechanisms through which that manipulation actually operates. Whether Google's interface design exploits behavioral biases, limited attention, default inertia, visual salience, or whether it simply reduces search costs for rational users, is not a pedantic theoretical question. It determines what remedies can work. And it determines whether the Digital Markets Act, the EU's regulatory framework for digital gatekeepers, is actually equipped to address the harm it was designed to prevent.

The behavioral economics literature offers a precise way into this problem. Heidhues, Köster and Kőszegi (2023) show that the welfare effects of steering depend fundamentally on who is being steered. Steering a rational consumer is broadly benign: she evaluates the recommendation and rejects it if it does not serve her interests. Steering a fallible consumer, one whose attention is limited, whose inertia is strong, and who may not even notice the manipulation, can generate welfare losses that are potentially unbounded. No structural remedy corrects this. The distinction matters because the existing regulatory apparatus, including the DMA, was designed with rational consumers in mind.

This dissertation investigates whether the behavioral biases literature can resolve the ambiguity at the center of the Google Shopping case. Two hypotheses organize the analysis. Under H1, Google's choice architecture, the prominent display of Shopping units, rich snippets, star ratings, and the systematic downranking of competitors, exploits cognitive biases to generate an advantage that persists independently of product quality. Under H2, the same design elements reduce consumer search costs, inform rational decision-making, and reflect genuine quality superiority rather than behavioral manipulation. The two hypotheses are observationally equivalent in many settings: users click on Google's results in both cases.

Distinguishing them requires either experimental variation in interface design at constant quality, or evidence about user awareness and post-click outcomes. Much of that evidence exists, but rarely in the form that would constitute a decisive test.

The approach is qualitative and theory-driven. It draws on the behavioral industrial organization literature, experimental economics, and EU competition law. Empirical evidence comes from field experiments (Allcott et al., 2025; Ursu, 2018), natural experiments exploiting regulatory variation (Decarolis et al., forthcoming), and the documentary record of Case AT.39740. No original data is collected. The contribution is analytical rather than empirical. The central finding is that the existing evidence is more consistent with H1 than with H2, particularly in high-inattention contexts such as mobile search, but that a decisive experimental test has not yet been conducted. The Digital Markets Act addresses the structural dimensions of self-preferencing but remains insufficiently calibrated for persistently inattentive consumers. No choice screen or ranking obligation can fully close that gap.

The analysis proceeds as follows. Chapter 1 establishes the economic and legal context of the Google Shopping case. Chapter 2 reviews the relevant academic literature, from industrial organization models of consumer search to the behavioral economics of steering and inattention. Chapter 3 identifies the three principal behavioral mechanisms through which self-preferencing may operate: limited attention, default bias, and visual salience. Chapter 4 develops the analytical framework, formalizing H1 and H2 as competing hypotheses and assessing the available empirical evidence. Chapter 5 applies this framework to the Google Shopping case directly. Chapters 6 and 7 turn to the Digital Markets Act, mapping its obligations against the behavioral literature and identifying what its provisions leave unaddressed. The conclusion draws together the findings and advances two concrete policy recommendations.

# Chapter 1 — Self-Preferencing in Digital Markets: Mechanisms, Evidence, and the Limits of Structural Remedies

## 1.1 The economics of gatekeeping platforms

The concept of self-preferencing did not originate with the Google Shopping case. It describes a broader pattern that emerges when a firm occupies two positions at once: it operates a marketplace through which independent providers reach consumers, and it also competes within that marketplace as a seller. The economic literature refers to such firms as "hybrid platforms" or "dual-role platforms" (Kittaka, Sato & Zennyo, 2023). The potential for harm comes from the informational and architectural asymmetries the platform controls. Because the platform designs the interface through which consumers discover and evaluate products, it can influence what is seen, how prominently, and in what order regardless of the underlying quality of the options available.

The structural argument against self-preferencing follows directly. A platform that ranks its own services more favourably than competing services of equivalent quality forecloses competition not through better products but through privileged distribution. Rivals cannot replicate that advantage because it is architectural, not quality-based. Entry is discouraged, incumbents are entrenched, and consumers face less competitive pressure on price and quality. This is the core theory of harm in Case AT.39740. It has since been reproduced, with

variations, in enforcement actions against Amazon's Buy Box, Apple's App Store, and Meta's advertising practices.

What distinguishes the digital context from earlier antitrust cases is the scale and opacity of the mechanism. Traditional cases involving shelf-space allocation or exclusive dealing involved visible, contractual arrangements between identifiable parties. In a search engine, the arrangement is encoded in an algorithm. The ranking criteria are proprietary. The effects on traffic are measurable in aggregate but difficult to attribute causally to any specific design choice. This opacity is not incidental. It is, as this chapter will argue, central to understanding why behavioral remedies may be necessary where structural ones are insufficient.

## 1.2 The Google Shopping case: a primer

Case AT.39740 was initiated by the European Commission in 2010 following complaints from rival comparison shopping services. The investigation focused on a specific pattern of conduct. Google's search results page presented its own comparison shopping service, initially known as Froogle, later as Google Shopping, in a visually prominent Shopping Unit positioned at or near the top of the page, accompanied by product images, prices, and star ratings. Competing services appeared only in the organic results, where they were subject to Google's standard ranking algorithms. The Commission found that these algorithms had been deliberately adjusted to demote competing services.

The quantitative evidence was substantial. Click-through rates for results in the first organic position are dramatically higher than for results lower on the page. The Commission's data showed that the Shopping Unit achieved click-through rates that rival services in positions four through ten could not approach, regardless of their quality scores. The Commission also documented what it called "search divergence": the systematic displacement of relevant organic results, including those from rival comparison services, below the visible fold of the page. Google's own internal studies, cited in the decision, acknowledged that position on the results page was a primary determinant of traffic.

The decision was upheld in all material respects by the General Court in 2021 and confirmed by the Court of Justice of the European Union in 2024. The legal foundation rests on Article 102 TFEU and the principle that a dominant firm abuses its position when it forecloses rivals through means other than competition on the merits. The Court rejected Google's argument that the Shopping Unit improved the quality of search results for users. Benefits to users, it found, could not justify the competitive foreclosure of rival services that were never given the opportunity to compete on equal terms.

Two aspects of the legal reasoning matter for this dissertation. First, the General Court rejected the counterfactual studies submitted by Google, which attempted to show that the visibility differential had only a marginal effect on traffic. The Court found that these studies examined the upgrading of Google Shopping in isolation, without considering the simultaneous downgrading of competing services. That methodological critique anticipates the H1/H2 distinction developed in Chapter 4. Second, and more fundamentally, neither the Commission nor the Court engaged with the question of why users responded to the visibility differential as they did. The assumption throughout was that users followed the results they were shown. The mechanism behind that behaviour, whether rational inference, limited attention, or visual salience, was treated as irrelevant to the legal analysis.

## 1.3 Forms of self-preferencing: a taxonomy

The Google Shopping case illustrates one form of self-preferencing, but it is not the only one. Kittaka, Sato and Zennyo (2023) identify two principal forms in their review of the literature. The first is the manipulation of search result ordering and ranking in favour of a platform's own services. The second is the exploitative use of third-party data, specifically using transaction and behavioural data generated by independent sellers on the platform to inform the platform's own product development and pricing decisions. Amazon's use of third-party seller data to identify and undercut successful products is the most well-known example of the second form. The DMA addresses both: Article 6(5) covers ranking, Article 6(2) covers data use.

This dissertation focuses on the first form: ranking manipulation and the design of the search results page. This choice is analytically motivated. The behavioral biases through which ranking manipulation operates, attention, salience and inertia, are directly observable in user behaviour and testable in experimental designs. The behavioral channels through which data exploitation harms consumers are more diffuse and operate primarily through market structure rather than through individual cognitive responses.

Within the category of ranking manipulation, three specific mechanisms appear in the Google Shopping case and will recur throughout the analysis. The first is prominence by elevation: the systematic placement of Google's own services in positions of high visual salience, typically at the top of the page and in formats, images, star ratings and prices, that attract disproportionate attention relative to plain-text organic results. The second is prominence by displacement: the demotion of competing services to positions where users are substantially less likely to encounter them, whether through explicit algorithmic adjustment or through the allocation of above-the-fold space to the platform's own content. The third is prominence by default: the pre-selection of Google's services as the default search engine on Android devices and in Chrome, which creates an initial condition of use that persists through inertia even when users would, if prompted, prefer an alternative.

These three mechanisms are conceptually distinct. Elevation operates through the salience and attention channels. Displacement operates through the search cost channel, increasing the cognitive effort required to reach competing services. Default operates through the inertia channel. The regulatory implications of each differ, and the current DMA framework does not fully reflect that distinction.

## 1.4 Why structural remedies may be insufficient

The standard response to self-preferencing is structural. Regulators require the platform to apply the same ranking rules to its own services as to third-party services. They prohibit preferential placement. They mandate that competing services appear in equivalent formats. Article 6(5) of the DMA follows this logic.

The approach makes sense under one assumption: that users, once shown competing options on equal terms, will compare them and pick the best one.

But that assumption is exactly what the behavioral economics literature questions. Many users do not actively compare their options. They are inattentive, inertial, or strongly influenced by

visual design. For these users, equal ranking conditions may not be enough. A user who never considers the alternatives cannot choose among them, no matter how prominently those alternatives appear. The harm is not that she was blocked from seeing a better option. It is that the platform's design exploited a cognitive limitation to make sure she never thought about it.

This is where standard IO reaches its limit. It assumes that users respond rationally to the information they receive. Remove the structural barrier, and competition follows. The behavioral IO literature challenges that conclusion. Heidhues, Köster and Kőszegi (2023) show why structural remedies may leave the core of the problem untouched. Chapter 2 reviews their framework and the broader behavioral literature, building the analytical tools needed to address the full complexity of the self-preferencing problem.

## Chapter 2 — From Rational Search to Behavioral Exploitation: A Review of the Literature

### 2.1 Search costs, position effects, and the rational baseline

Why do markets with many competing sellers not always produce competitive prices? Stigler (1961) gave the first rigorous answer: gathering information is costly. Consumers do not compare every available option. They search until the expected gain from one more search falls below its cost, then stop. The result is price dispersion. Firms facing less attentive consumers can charge more without losing them. Frictions, not market power in the traditional sense, drive the distortion.

Online platforms seemed, at first, to promise a solution. If comparing products across hundreds of sellers takes seconds, consumers should search more, prices should converge, and competition should intensify. Early evidence was broadly consistent with this expectation. But the same literature quickly identified a problem. Platforms do not simply reduce search costs: they shape them. By controlling which products appear, in what order, and in what format, a platform decides which alternatives consumers are likely to consider. The search results page is not a neutral display. It is an economic choice.

The directed search literature formalises this point. Haan, Moraga-González and Petrikaitė (2018) model consumers who observe seller signals and decide which options to evaluate. Their central result is that when a platform grants one seller disproportionate prominence, even rational consumers visit the disadvantaged seller less often, regardless of its quality. No cognitive bias is needed. Prominence alone redirects attention. Liu and Dukes (2016) reach a similar conclusion. When consumers can only evaluate a subset of products, the composition of that subset depends partly on platform design. If prominence reflects commercial affiliation rather than quality, the subset is systematically distorted. De los Santos, Hortaçsu and Wildenbeest (2012) confirm empirically that search intensity responds to cost structures in ways broadly consistent with rational search models.

These contributions establish an important baseline. Even if Google's users are fully rational, platform design choices that favour affiliated services can produce anticompetitive outcomes. The question this dissertation investigates is whether behavioral mechanisms make that harm worse. Two experimental studies speak directly to this point.

Ursu (2018) provides the most rigorous causal estimate of position effects in an online search context. Using data from a field experiment at Expedia where product rankings were randomly assigned, she finds that position has a large causal effect on which products consumers click. Higher-ranked products receive substantially more clicks. But conditional on clicking, position has no significant effect on purchase decisions. Consumers use position to decide what to look at, not what to buy. Position effects work through the search cost channel: a higher-ranked product is cheaper to reach, not intrinsically better. This is efficient if position reflects quality. The welfare concern arises when it does not.

Dinerstein, Einav, Levin and Sundaresan (2018) document a related tradeoff at eBay. A change in platform design improved matching between consumers and products but weakened price competition by concentrating demand on fewer sellers. Better matching and more concentrated demand are not independent outcomes. When a platform steers consumers toward a preferred subset of options, it reduces competitive pressure on the sellers in that subset.

Both studies assume rational consumers. The following sections examine what changes when that assumption is relaxed.

## 2.2 The personalization paradox: Kaye (2025) and the pricing response problem

Kaye (2025) asks a simple but important question: what happens to consumer welfare when a platform improves its recommendation system? The standard answer is that better recommendations help consumers. They find products that match their preferences faster, search costs fall, and surplus rises. Kaye shows this answer is incomplete. It ignores how sellers respond.

The paper uses clickstream data from Expedia Group, covering over 330 000 consumer searches across 136 000 hotel properties. The setting is well suited to the question. Expedia controls its recommendation system but third-party hotels set prices independently. This structure, platform chooses recommendations and sellers choose prices, mirrors the architecture of Google Shopping exactly.

Kaye develops a structural model of demand, platform recommendations, and hotel pricing. The counterfactual simulations compare four levels of personalisation, from a basic system using only product features to a fully personalised system using past purchases and browsing data.

The results are striking. When prices are held fixed, personalisation increases consumer surplus by approximately 2.3% of total booking revenue, around 0.9 billion dollars at Expedia's 2013 scale. Better matching helps consumers find suitable products faster. This is the long-tail effect.

But once sellers can update prices, the result reverses. Hotels facing more precisely targeted, less price-sensitive consumers raise prices. Quantity falls by 4.5%. Hotel profits rise by 4.9%. Consumer surplus declines by 5% of booking revenue, approximately 2 billion dollars. The welfare loss is nearly double the initial gain.

The mechanism is simple. Better recommendations concentrate demand. Sellers facing a more captive consumer base respond by charging more. The efficiency gain from matching does not offset the loss from weaker price competition.

The parallel with Google Shopping is direct. When Google promotes its own services in prominent positions, it concentrates consumer attention on a specific set of providers. Those providers face a more captive, less price-sensitive demand than they would under a neutral ranking system. The harm is not just that competitors receive less traffic. It is that concentrated demand enables above-competitive pricing, with losses borne by consumers who were, in principle, directed toward a better match.

Kaye's findings also point to the limits of purely structural remedies. Requiring Google to rank competing services equally addresses the allocation of prominence but not the pricing incentives that concentrated demand creates. Even a neutral ranking system may generate demand concentration if inattentive users consistently treat Google's own services as the default. This connects directly to the behavioral framework developed in the next section.

## 2.3 Behavioral Industrial Organization: the theoretical core

The models reviewed in section 2.1 establish that platform design can distort competition even when users are fully rational. But they also imply a natural limit. A rational consumer who recognises that a prominent result reflects commercial affiliation rather than quality can adjust. She scrolls further down, searches elsewhere, or simply ignores the nudge. The harm is real but bounded.

The Behavioral Industrial Organization literature challenges that limit. Its central claim is that a substantial fraction of consumers do not behave as the rational model assumes. They are not irrational in a general sense. But their attention is limited, their habits are persistent, and careful deliberation is often not worth the effort in routine, low-stakes decisions. Firms that understand these patterns have strong incentives to exploit them, and in equilibrium, they do.

The foundational framework comes from Heidhues and Kőszegi (2018), whose handbook chapter on Behavioral IO establishes the tools used throughout the subsequent literature. Their key insight is that the presence of naive or inattentive consumers changes the competitive equilibrium in ways that markets cannot self-correct. Sophisticated firms learn which consumers are susceptible to which biases and design their strategies accordingly. Naive consumers are harmed directly. Sophisticated consumers may be harmed indirectly, because firms optimise for the population as a whole rather than for each type separately. In a search engine context, the relevant form of naivety is not general cognitive impairment. It is the tendency to follow the first salient result without evaluating alternatives.

The most directly relevant contribution for this dissertation is Heidhues, Köster and Kőszegi (2023). Their paper analyses the welfare effects of steering, defined as any platform design choice that influences which product a consumer considers first, under two different assumptions about consumer rationality. The distinction they draw is fundamental.

Let U* denote the first-best utility achievable under neutral platform design. When a rational consumer is steered toward a product, she evaluates it and adopts it only if it is genuinely better than the alternatives. The welfare loss is therefore small and bounded. She retains the ability to correct for the steering once she engages with the choice.

When a fallible consumer is steered, the calculus is fundamentally different. A fallible consumer does not evaluate alternatives. She adopts the steered product without comparison, regardless of whether it serves her interests. The platform can therefore extract unlimited surplus from her by consistently directing her toward high-margin products that she would not have chosen under full information. The welfare loss is not bounded. It may be arbitrarily large.

The distinction between rational and fallible consumers captures the theoretical motivation for this dissertation's central question. If Google's users are predominantly rational, recognising that prominent results reflect commercial affiliation and adjusting accordingly, the welfare effects of self-preferencing are modest and likely self-correcting. Structural remedies that equalise the conditions of prominence may be sufficient. But if a substantial fraction of users are fallible in the sense of Heidhues et al. (2023), persistently inattentive, strongly inertial, or heavily influenced by visual design, the welfare loss is potentially large and not self-correcting. Structural remedies that ignore the behavioral mechanism will fail.

The paper also distinguishes between types of steering signals. An informative signal conveys genuine quality information, something true about the product that the consumer would value knowing. An uninformative signal attracts attention without conveying relevant content. A misleading signal conveys false quality information. The welfare implications differ. Informative steering of a rational consumer can improve welfare. Any steering of a fallible consumer is potentially harmful, regardless of whether the signal is informative, because the fallible consumer cannot assess its accuracy.

This typology connects directly to the H1/H2 distinction developed in Chapter 4. Under H2, Google's visual elements, rich snippets, star ratings and high-resolution images, are informative signals. They help rational consumers find better products faster. Under H1, they are uninformative or misleading signals that exploit visual salience to capture the attention of fallible consumers, directing them toward Google's affiliated services regardless of comparative quality.

One implication follows from this framework for the rest of the dissertation. Whether self-preferencing causes significant and lasting harm does not depend on whether prominence effects exist. They do. It depends on the proportion of consumers who are fallible. Chapter 3 turns to the empirical evidence on that proportion, examining what the experimental literature reveals about limited attention, default bias, and visual salience in online search environments.

## Chapter 3 — Behavioral Mechanisms: How Self-Preferencing Operates in Practice

This chapter identifies the three behavioral mechanisms through which Google's self-preferencing may operate: limited attention, default bias, and visual salience. For each mechanism, the analysis follows the same structure. It first establishes the theoretical channel, then reviews the experimental evidence, and finally applies the findings to the specific context of Google Search. The chapter does not resolve whether Google exploits these mechanisms or merely informs rational consumers — that question is addressed in Chapter 4. The purpose here is to establish that the mechanisms exist, that they are empirically documented, and that they are plausibly active in the Google Shopping context.

## 3.1 Limited attention

Standard economic models assume that consumers consider all available options before making a choice. In practice, attention is scarce. Consumers do not evaluate every result on a search page. They focus on what is immediately visible, click on what seems most relevant at first glance, and rarely look further. This is not irrational behaviour. It is a sensible response to the cost of deliberation when the stakes are low and the options are many. The problem arises when platforms are designed in ways that systematically direct that limited attention toward affiliated services rather than the most relevant alternatives.

The most direct evidence on limited attention in online search comes from Allcott, Castillo, Gentzkow, Musolff and Salz (2025). Their field experiment recruited a large sample of US internet users and assigned them to different conditions that modified their default search engine or required them to make an active choice. The experiment was designed to decompose the sources of Google's market power into three components: genuine quality advantage, structural switching costs, and persistent inattention.

The key finding for this section is that approximately one third of users in the sample exhibit persistent inattention. These users do not actively consider their choice of search engine. They use whatever is set as the default, and they continue to do so even when the alternative performs comparably. Crucially, when these users are switched to Bing by default, many of them do not return to Google, not because they prefer Bing, but because they never actively evaluate the choice in either direction. They follow the default passively, regardless of which engine it is.

This result has a direct implication for the Google Shopping case. If one third of users never actively consider which search engine they use, it is reasonable to assume that a similar fraction never actively consider which comparison shopping service they consult. When Google Shopping appears prominently at the top of the results page, these users are likely to click on it not because they have evaluated it against competing services, but because it is the first salient option they encounter. Their attention is captured before deliberation begins.

Decarolis, Li and Paternollo (forthcoming) provide complementary evidence from a different angle. They study the effect of regulatory interventions that changed the default search engine on Android devices in the European Economic Area, Russia, and Turkey. In the EEA, where the intervention required device manufacturers to present users with a choice screen at setup, Google's market share fell by less than one percentage point. In Russia and Turkey, where a competing engine, Yandex, was either mandated or strongly promoted as the default, Google's market share fell by more than ten percentage points.

The contrast between these two outcomes is informative. A choice screen presents users with options but requires them to make an active decision. A default switch simply changes what users receive without requiring any decision at all. The fact that a default switch produces an effect ten times larger than a choice screen suggests that the binding constraint on user behaviour is not the availability of alternatives. Users in all three jurisdictions had access to competing search engines. The constraint is the willingness to engage in active deliberation. Users who are given a choice screen must decide. Many choose not to, and stay with the pre-selected option. Users who are given a different default simply follow it.

Information about the quality of competing search engines does not appear to change this pattern. Allcott et al. (2025) find that users who are forcibly exposed to Bing update their beliefs about its quality upward. But this belief update does not translate into sustained switching. The inattentive third of users who were switched to Bing by default did not return to Google, but they also did not make an active choice to stay with Bing. They simply remained passive. The barrier to switching is not a lack of information about alternatives. It is the absence of active deliberation, and information alone cannot overcome that.

## 3.2 Default bias and the limits of choice architecture

The evidence reviewed in section 3.1 identifies a population of persistently inattentive users who follow defaults without deliberation. But inattention and default bias are not identical concepts, and the distinction matters for policy. A consumer who sticks with a default because the switching cost is high, because changing settings takes time, requires technical knowledge, or involves learning a new interface, is responding rationally to a structural friction. A consumer who sticks with a default because she never considered whether to switch is exhibiting a behavioral bias. The two patterns look identical in observational data. Both produce inertia. But they have radically different implications for what remedies can work.

If the source of inertia is structural switching costs, reducing those costs should be sufficient to restore competitive outcomes. A well-designed choice screen that makes switching easy and frictionless should produce significant reallocation of users across competing services. If the source of inertia is persistent inattention, reducing switching costs is not enough. A consumer who does not consider the choice will not respond to a choice screen, regardless of how easy switching has been made. What would be required is not a reduction in switching costs but a forced exposure to the decision, something that actively interrupts the user's passive behaviour and requires engagement.

Allcott et al. (2025) provide a clean empirical test of which mechanism dominates. Their experimental design includes a condition in which users are presented with an active choice between Google and Bing before their search session begins. This condition eliminates structural switching costs entirely: the user is presented with a simple interface, the options are clearly labelled, and selecting Bing requires one click. If switching costs were the primary source of inertia, this condition should produce a large increase in Bing's market share. The actual increase is 1.1 percentage points. This is a precisely estimated and statistically significant effect, but it is very small. It implies that structural switching costs account for only a minor share of Google's market power.

The implication is direct. The dominant source of default inertia in online search is persistent inattention, not structural switching costs. Users who do not consider the choice will not respond to a choice screen, because a choice screen only helps users who are willing to engage with the decision it presents. For the inattentive third of the user population, the relevant intervention is not a reduction in friction but an interruption of passivity, and no current regulatory instrument reliably achieves that.

This finding has a precise application to the Google Shopping context. Google's status as the default search engine on Android devices and in the Chrome browser creates an initial condition of use that persists through inattention. When a user opens Chrome and types a query, she is directed to Google Search by default, and from there to Google Shopping results

that appear prominently at the top of the page. At no point in this sequence does she make an active choice. The default structure removes the need for deliberation, and for inattentive users, the absence of deliberation is sufficient to generate and sustain Google's market position.

The evidence from Decarolis et al. (forthcoming) reinforces this conclusion. The EEA choice screen intervention, which reduced structural switching costs to near zero for users who engaged with it, produced a market share effect of less than one percentage point. The Russian and Turkish default-change interventions, which did not require any active engagement, produced effects ten times larger. The most effective way to change the behaviour of inattentive users is not to offer them a choice. It is to change what they receive without requiring a choice.

## 3.3 Visual salience and the design of the search results page

The two mechanisms reviewed above, limited attention and default bias, operate primarily at the level of search engine choice. They explain why users end up on Google rather than Bing or DuckDuckGo. But self-preferencing in the Google Shopping case operates at a different level: within the Google search results page, where Google's own comparison shopping service competes with third-party services for the user's attention. Understanding this within-page mechanism requires a different theoretical framework, one that addresses how visual design shapes attention allocation among options that are simultaneously visible on the same page.

Bordalo, Gennaioli and Shleifer (2013) provide the foundational theory. Their model of salience and consumer choice starts from a simple observation: attention is not allocated uniformly across the attributes of available options. Some attributes stand out and receive disproportionate weight in the decision process. An attribute is salient when it differs substantially from the average value of that attribute across the choice set. A price that is much lower than average, a rating that is much higher than average, or an image that is much larger than those of competing products will attract disproportionate attention relative to their actual informational content.

The welfare implications of this mechanism depend on whether the salience is warranted. If a product's salient attribute reflects genuine quality superiority, then salience directs attention efficiently and can reduce search costs. If the salience is manufactured, if the product achieves visual prominence through design choices that are independent of quality, then salience misleads. It directs attention toward an option that is not necessarily better, at the expense of alternatives that may be more suitable.

Li and Camerer (2022) provide experimental evidence that salience effects operate in precisely the way Bordalo et al. (2013) predict. In a series of controlled experiments, they show that visual salience influences economic choices in a predictable and automatic way. Importantly, these effects are not eliminated by deliberation. Even when subjects are aware that the choice environment contains a salient stimulus, they continue to over-weight it relative to less prominent alternatives. Salience effects are not a product of carelessness or ignorance. They reflect a fundamental feature of how human attention operates, one that persists even when the decision-maker is aware of it.

Applied to the Google search results page, this framework generates a precise prediction. Google Shopping units typically appear at the top of the results page in a visually distinctive format: high-resolution product images, star ratings displayed in gold, prices listed prominently, and a horizontal carousel layout that is visually different from the plain-text organic results below. Each of these design elements increases the visual salience of Google Shopping relative to competing services. According to Bordalo et al. (2013), users will allocate disproportionate attention to Google Shopping, not necessarily because it is better, but because it is more visually prominent.

The critical question, which this dissertation cannot resolve with the currently available evidence, is whether this salience is informative or manufactured. Under H2, the rich snippets of Google Shopping convey genuine quality information: product images help consumers assess fit, star ratings summarise user reviews, and price displays reduce the cost of comparison. Under H1, the visual design of the Shopping unit is engineered to capture attention regardless of comparative quality. The salience is a product of design choices that Google controls, not a reflection of the underlying quality of the search results.

Distinguishing between these two interpretations requires evidence on whether Google Shopping's click advantage persists after controlling for the quality of the results it returns. As this dissertation will argue in Chapter 4, the existing literature does not yet provide that evidence in the form of a clean experimental test. What the evidence does establish is that the mechanism exists. Visual salience systematically directs attention toward prominent options, the effect persists even among deliberate decision-makers, and Google's interface design creates precisely the conditions under which salience effects are largest.

Three conclusions follow from the evidence reviewed in this chapter. First, a substantial fraction of Google's users, approximately one third by the estimate of Allcott et al. (2025), never actively consider their choice of search engine. For these users, the default determines the outcome, and no amount of information about competing alternatives changes that. Second, the dominant source of this inertia is persistent inattention rather than structural switching costs. Reducing friction is not sufficient to restore active deliberation among users who do not deliberate. Third, the visual design of the Google search results page creates conditions of high salience that systematically direct attention toward Google Shopping before rational evaluation can occur. Whether that salience is informative or exploitative is the question to which Chapter 4 now turns.

# Chapter 4 — Exploitation or Information? An Analytical Framework

## 4.1 The central dilemma

The previous chapters established two things. First, behavioral mechanisms, limited attention, default bias, and visual salience, are empirically documented and plausibly active in the Google Search context. Second, competition law has focused on the structural asymmetry between Google Shopping and rival comparison services without asking how users actually process the information they are shown. The Commission proved that Google diverted traffic. It did not explain why users followed.

This chapter addresses that gap directly. It asks whether users click on Google Shopping because it is genuinely better, in which case the prominence of Google Shopping reflects quality, or because the interface is designed in a way that captures attention before users can

evaluate alternatives. These two possibilities are formalised as H1 and H2. The rest of this chapter derives their empirical predictions, evaluates the available evidence, and identifies what a decisive test would require.

The distinction matters beyond academic interest. Under H1, equalising the structural conditions of competition, requiring Google to rank competing services on the same terms as its own, may not be enough. If the harm operates through visual salience and persistent inattention, a neutral ranking rule leaves the behavioral mechanism intact. Under H2, structural remedies are sufficient. Getting the diagnosis right is a precondition for getting the remedy right.

## 4.2 Formalising the two hypotheses

The theoretical framework draws on Heidhues, Köster and Kőszegi (2023) and the salience model of Bordalo, Gennaioli and Shleifer (2013). The goal is not a fully specified model. It is to give H1 and H2 enough structure to derive testable predictions.

### *The rational consumer baseline — H2*

Under H2, the consumer is fully rational. She uses the visual design of the results page, images, star ratings, position, as a signal of underlying quality. A high-ranked, visually rich result is treated as more likely to be relevant, and she is correct to do so if the ranking reflects quality. The interface reduces search costs without distorting her final choice. It is welfare-improving.

### *The behavioral exploitation hypothesis — H1*

Under H1, perceived utility differs from true utility. The visual design introduces a salience premium: a result with more striking attributes, larger images, gold star ratings, prominent prices, appears more attractive than it would if those attributes were stripped away. Following Bordalo et al. (2013), this premium is larger when a result stands out relative to the rest of the choice set. Google Shopping results systematically stand out more than competing plain-text organic results. If the salience premium is large enough, users click on Google Shopping even when its underlying quality is no higher than that of competing results.

This premium might seem easily corrected. A user who is aware that visual design influences her attention could, in principle, deliberately discount it and evaluate results on content alone. If this were true, H1 would be self-limiting: sophisticated users would learn to ignore the visual design, and the mechanism would only affect genuinely inattentive users.

Li and Camerer (2022) show that this correction does not reliably occur. In a series of controlled experiments, they find that visual salience influences economic choices in a predictable and automatic way, what psychologists call a bottom-up attentional effect, driven by the stimulus itself rather than by the decision-maker's goals. Critically, this effect persists even when subjects are explicitly aware that the choice environment contains a salient stimulus and are instructed to evaluate options on their merits. Deliberation reduces the salience effect but does not eliminate it.

The implication is direct. Under H2, a user who recognises that Google Shopping's visual prominence reflects commercial design can adjust her behaviour and choose the best result regardless of format. Under H1, even a user who suspects that prominence reflects commercial interest cannot fully suppress the attentional capture that visual salience generates. The salience premium remains positive. This means that the population susceptible to the H1 mechanism is larger than the population of persistently inattentive users identified by Allcott et al. (2025). It includes not only the third of users who never actively consider their options, but also a fraction of deliberating users who still over-weight the salient result.

The welfare loss under H1 therefore depends on the share of fallible users in the population. As established in Chapter 2, the loss for rational users is small and bounded while the loss for fallible users may be arbitrarily large. For a share of fallible users around one third, the estimate from Allcott et al. (2025), and a large welfare loss per fallible user, the aggregate harm can be substantial. Whether H1 or H2 better describes the world is therefore partly an empirical question about how many users are fallible and partly an empirical question about how large the harm per fallible user actually is.

## 4.3 The observational equivalence problem

The core difficulty in distinguishing H1 from H2 is that both hypotheses predict the same surface outcome: users click on Google Shopping at high rates. This is what economists call observational equivalence. The data generated under H1 and the data generated under H2 look identical in a standard dataset that records only clicks and purchases.

Consider the first position on a search results page. Across multiple studies, the first organic result captures approximately 28 to 34 percent of clicks, and this share drops steeply for lower positions. This pattern is consistent with H2: rational users click first on the most prominent result because they believe it is the most relevant. But the same pattern is equally consistent with H1: users with limited attention click on whatever is most salient, and the first result is most salient by construction.

The same problem applies to conversion rates. If Google Shopping users buy at similar rates to users who found comparable products through other channels, this is consistent with H2, the product was genuinely good. But it is also consistent with H1, the user was directed to an adequate product, just not necessarily the best available. Without a counterfactual, what would the user have bought if she had not been steered toward Google Shopping, conversion rates cannot distinguish between the two hypotheses.

This is why click data alone cannot resolve the H1/H2 question. The Commission documented that users click on Google Shopping at high rates. It did not establish whether those clicks reflect quality or capture. Resolving the question requires experimental designs that break the observational equivalence, either by holding quality constant while varying visual design, or by holding visual design constant while varying quality.

## 4.4 What the existing evidence shows

The experimental literature reviewed in Chapters 2 and 3 provides partial but not decisive evidence on H1 versus H2. Three studies speak most directly to the question.

Ursu (2018) is the most methodologically rigorous study of position effects in an online search context. Her key finding is that position affects which results users click but not which results they ultimately purchase, conditional on having searched. If position were a pure quality signal under H2, it should affect both clicks and purchases. Users who trust the ranking should both click on higher-ranked results and buy from them more often. The fact that position affects clicks but not purchases suggests that users update on true quality once they search. Their initial click was therefore driven by something other than quality, which is consistent with a search cost or salience interpretation rather than a pure H2 mechanism.

Allcott et al. (2025) provide the most direct evidence on the composition of the user population. Their estimate that approximately one third of users are persistently inattentive sets a floor on the share of fallible users in the welfare loss established in Chapter 2. For this third of users, H2 cannot hold by definition: users who never actively consider their choice of search engine cannot be responding to quality signals. Their behaviour is driven entirely by the default and the interface design they encounter. This does not prove H1 for the full population, but it establishes that H2 cannot account for the full pattern of user behaviour.

A further argument against H2 is that Google's quality advantage may not be intrinsic to its algorithm. It may simply reflect the fact that Google has more users, and therefore more data to train its ranking system. As noted by Allcott et al. (2025), citing testimony from United States v. Google LLC (2023), Microsoft executives confirmed that further investment in mobile search would not be economically viable without access to Google's scale of data. If Bing had access to the same volume of search data, it might perform equally well. Under this interpretation, Google's apparent quality superiority is itself a product of its market power, not an independent justification for it.

Kaye (2025) adds a further dimension. His finding that improved recommendation systems reduce consumer surplus once seller pricing responses are accounted for shows that the efficiency argument for platform design, the H2 claim that better interface design helps consumers, is incomplete even on its own terms. A platform that improves matching also concentrates demand, and concentrated demand enables above-competitive pricing. The net welfare effect of a seemingly pro-consumer design choice can be negative. This does not directly test H1 versus H2, but it weakens the H2 defence of self-preferencing: even if Google Shopping's visual design genuinely helps users find better products, the concentration of demand it creates may still harm them through higher prices.

## 4.5 The central gap and proposed experimental designs

The analysis above leads to a clear finding. The existing evidence is more consistent with H1 than H2, particularly the evidence that persistent inattention, not quality or switching costs, is the dominant source of Google's market power. But none of the existing studies constitutes a decisive test. The reason is that no published study has simultaneously held product quality constant and varied the visual design of search results in the specific context of Google Shopping.

This is the central gap in the literature. Without it, the question of whether Google Shopping's click advantage reflects the informational value of its visual design or the attentional capture that salience generates independently of content cannot be answered. The following experimental designs would provide the necessary evidence.

Experiment 1 — Visual ablation. Users would be randomly assigned to see Google Shopping results in their standard rich snippet format, images, star ratings and prices, or in a plain-text format identical to that of organic results. The content of the results would be held constant across conditions: same products, same prices, same relevance scores. The only thing that varies is the visual format. If the click advantage of Google Shopping disappears in the plain-text condition, this confirms that the advantage is driven by visual salience rather than content quality, consistent with H1. If the click advantage persists in both conditions, the result supports H2.

Experiment 2 — Rich snippet randomisation. Users would be randomly assigned to conditions in which the rich snippet format, images, star ratings and price displays, is assigned either to Google Shopping or to a competing comparison service, while the content of the results and their position on the page are held constant. This design isolates the effect of the visual format specifically, independently of position and content. If click rates follow the rich snippet format, users click on whichever service has the images and stars regardless of which it is, this confirms that visual salience drives clicks independently of quality, supporting H1. If click rates do not change, the result supports H2. This experiment is directly relevant to the interpretation of Article 6(5) of the DMA. If the click advantage of Google Shopping is driven by its rich snippet format rather than by its ranking position, then a ranking neutrality obligation alone would not eliminate the behavioral source of the harm.

Experiment 3 — Post-click outcome measurement. Users who click on Google Shopping results would be tracked to measure post-click outcomes: time spent on the destination page, return rates, repeat purchase rates, and self-reported satisfaction. The same outcomes would be measured for users who click on competing comparison services for similar queries. Under H2, Google Shopping should deliver systematically better post-click outcomes. Under H1, there should be no systematic difference in post-click outcomes, despite the large difference in initial click rates. A finding of no quality advantage in outcomes, combined with a large click advantage, would be the clearest possible evidence for H1.

Experiment 4 — Informational alerts. Users would be randomly assigned to see a brief disclosure before their search, informing them that the first results are ranked partly on the basis of commercial affiliation rather than relevance alone, or to see no disclosure. If H1 is correct, the disclosure should reduce clicks on Google Shopping: users who understand that prominence reflects commercial interest rather than quality should discount it accordingly. If H2 is correct, the disclosure should have no effect: rational users already account for commercial influence and choose Google Shopping because it is genuinely better. The same logic applies to the commercial search context, though the experiment has not yet been conducted there.

These experiments remain, to date, absent from the published literature on Google Shopping specifically. The Digital Markets Act creates a regulatory context in which they could be required. Article 6(5) currently mandates non-discriminatory ranking conditions. It could be extended to require that Google demonstrate the informational value of its rich snippet formats through independent empirical testing, a requirement that would effectively mandate something close to Experiment 1. The absence of such a requirement is one of the specific limits of the current DMA framework, examined in detail in Chapter 7.

## 4.6 What this chapter establishes

This chapter has done three things. It has formalised H1 and H2 in terms precise enough to derive testable predictions. It has shown that the existing evidence, while more consistent with H1 than H2, does not constitute a decisive test. This is partly because of the observational equivalence problem and partly because no study has yet combined experimental variation in visual design with a quality control. And it has proposed four experimental designs that would close that gap.

The contribution is not a definitive answer to the H1/H2 question. It is a precise identification of why that question has not yet been answered, and what answering it would require. That finding has a direct implication for regulation. If current interventions, including Article 6(5) of the DMA, rest on an empirical foundation that is incomplete, their effectiveness cannot be properly evaluated. Chapter 7 develops this implication in the context of the Digital Markets Act.

# Chapter 5 — The Google Shopping Case Through a Behavioral Lens

The previous chapters developed the theoretical and empirical tools needed to analyse self-preferencing as a behavioral phenomenon. This chapter applies those tools to the Google Shopping case directly. The goal is not to retell the legal story, which was covered in Chapter 1, but to read the Commission's findings through the H1/H2 framework and show what that reading reveals.

## 5.1 The click-through rate data as behavioral evidence

The Commission's decision in AT.39740 documented a striking pattern. When Google Shopping units appeared on the results page, they attracted approximately 42% of total clicks. Competing comparison services, appearing in organic results below the Shopping unit, received a small fraction of that traffic. The Commission treated this as evidence of traffic diversion, a structural harm to competitors. That interpretation is legally correct. But it is incomplete as an economic account of what was happening.

Read through the H1/H2 framework, the 42% figure takes on a different meaning. Under H2, it reflects quality: users clicked on Google Shopping because it was genuinely the most relevant result. The visual prominence of the Shopping unit was simply an efficient way of surfacing the best option. Under H1, the figure reflects behavioral capture: users clicked on Google Shopping because the rich snippet format captured their attention before they could evaluate alternatives. The click share reflects the design of the results page, not the quality of the results.

The Commission's data cannot on their own resolve this ambiguity. A 42% click share is consistent with both hypotheses. What makes the figure more informative is the Commission's additional finding that competing services were systematically downranked regardless of their quality scores. If Google Shopping's click advantage reflected genuine quality superiority, there would be no reason to simultaneously demote high-quality competing services. The combination of upranking Google Shopping and downranking rivals is consistent with H1: the goal was to capture attention and redirect traffic, not to surface the best results for users.

Heidhues, Köster and Kőszegi (2023) illustrate this mechanism with a documented case from Bing: a simple A/B test that moved information from the body of an advertisement to its title increased annual revenue by more than 100 million dollars, without changing the content or quality of the results (Kohavi, Tang and Xu, 2020). The product was identical. The information was identical. Only the visual format changed, and the revenue effect was enormous.

This interpretation is reinforced by what the Commission documented about the visual design of the Shopping unit. Google Shopping results appeared in a format, with high-resolution images, star ratings, and price displays, that was visually distinct from the plain-text organic results below. The Commission used eye-tracking studies to show that this format attracted disproportionate user attention and generated significantly higher click-through rates than plain-text results at equivalent positions. This is precisely the salience mechanism described by Bordalo, Gennaioli and Shleifer (2013): attributes that stand out relative to the choice set receive disproportionate attention, regardless of whether that attention is warranted by underlying quality.

The Commission did not frame its findings in behavioral economics terms. It documented the salience effect empirically, through eye-tracking data and click-through rates, without connecting it to the theoretical literature on cognitive exploitation. That connection is what this dissertation provides. The Commission's data are not just evidence of structural harm. They are evidence that the salience mechanism of H1 was operating at scale.

## 5.2 The rejection of Google's counterfactual studies

Google submitted several counterfactual studies during the Commission's investigation, arguing that the visibility differential had only a marginal effect on user behaviour. These studies attempted to show that users clicked on competing services at rates comparable to what would have been expected under a neutral interface. The Commission rejected these studies. The General Court confirmed that rejection in 2021, and the Court of Justice upheld it in 2024.

The legal grounds for rejection were procedural and methodological. But the underlying economic reason is directly relevant to the H1/H2 framework. Google's studies focused exclusively on the upgrading element of its conduct, the promotion of Google Shopping, without accounting for the simultaneous downgrading of competing services. This is a critical omission. The behavioral harm does not operate only through the salience of Google Shopping. It also operates through the reduced visibility of competitors, which increases the effective search cost of reaching them.

A user whose attention is captured by a salient result at the top of the page does not simply choose between Google Shopping and an equally visible alternative. She chooses between Google Shopping and an alternative that requires additional scrolling, reduced salience, and a higher cognitive cost to reach. Google's studies asked whether users clicked on competing services when those services appeared on the page. They did not ask whether users would have clicked on competing services if those services had appeared in an equally prominent format. That is the relevant counterfactual under H1, and it is the one that the existing studies never tested.

The Court's rejection of Google's evidence is therefore coherent with the behavioral framework, even though the Court did not invoke behavioral economics explicitly. By insisting that the analysis must account for both the promotion of Google Shopping and the demotion of rivals, the Court implicitly required a framework in which the relevant harm is the full choice architecture of the results page. That is precisely what H1 predicts: the harm is not that Google Shopping is prominent, but that the combination of its prominence and its competitors' reduced visibility systematically redirects attention in a way that users cannot easily correct.

### 5.3 What the case establishes and what it leaves open

The Google Shopping case establishes three things that are relevant to this dissertation. First, Google's conduct had a measurable effect on user behaviour: the click-through rate data show that the design of the results page redirected substantial traffic toward Google Shopping and away from competing services. Second, that effect cannot be explained by quality alone: the simultaneous downranking of high-quality competing services is inconsistent with a pure H2 account. Third, the visual design of the Shopping unit, documented through eye-tracking studies in the Commission's decision, operated through the salience channel identified in the behavioral economics literature.

What the case does not establish is the precise mechanism through which the behavioral harm operated. The Commission documented that users followed the design of the results page. It did not determine whether they did so because of limited attention, default inertia, visual salience, or some combination of the three. It did not measure post-click satisfaction or compare user outcomes under the biased interface with outcomes under a neutral one. And it did not propose a remedy that addressed the behavioral source of the harm, only one that addressed the structural asymmetry in ranking.

This last point connects directly to the analysis of Chapters 6 and 7. The remedy imposed by the Commission, and later reinforced by the Digital Markets Act, requires Google to apply non-discriminatory ranking conditions. It does not require Google to demonstrate that its visual design is informative rather than exploitative. It does not mandate the experimental tests proposed in Chapter 4. And it does not address the persistent inattention of the third of users who, as Allcott et al. (2025) document, never actively consider their options regardless of how those options are presented.

The Google Shopping case is, in this sense, a landmark that points beyond itself. It proved that self-preferencing causes competitive harm. It did not prove that the harm is behavioral in origin and therefore requires behavioral remedies. That argument is what this dissertation makes. The following chapters examine whether the Digital Markets Act is designed to address a behavioral problem or only a structural one.

## Chapter 6 — The Digital Markets Act: Relevant Obligations and Behavioral Foundations

The Google Shopping case took seven years to investigate. It took four more years to reach a final judgment. And by the time the Court of Justice confirmed the Commission's decision in

2024, the market had already changed significantly. This experience made one thing clear: ex post competition law is too slow for digital markets. The Digital Markets Act was designed to address this problem. Instead of waiting for harm to occur and then investigating, the DMA imposes obligations on large platforms in advance. They must follow certain rules before any specific abuse is proven.

This chapter describes the DMA obligations most relevant to self-preferencing. It also shows that some of those obligations were shaped by economists who work in the exact tradition this dissertation draws on. The goal is not to evaluate whether the DMA is sufficient, that comes in Chapter 7. The goal here is to understand what it actually requires, and why it represents a genuine step forward.

## 6.1 Gatekeepers and the logic of ex ante regulation

The DMA applies to platforms designated as "gatekeepers." To qualify, a platform must provide core digital services and meet specific thresholds on revenue, number of users, and geographic presence. Google was designated as a gatekeeper in September 2023. The designation covers its search engine, its browser, its operating system, and its advertising services. Once designated, the obligations apply automatically and no specific abuse needs to be proven.

This matters for the behavioral argument of this dissertation. Ex post competition law, as applied in AT.39740, asks whether a specific conduct caused specific harm to specific competitors. That framing struggles with behavioral harm. When harm operates through the cumulative effect of interface design on millions of inattentive users, there is no single act to point to. The Commission's investigation lasted seven years partly for this reason. The DMA sidesteps that problem by prohibiting certain design choices outright, regardless of their proven effect in any specific case.

This is, implicitly, an acknowledgment that some harms are structural. From the perspective of this dissertation, it is also an acknowledgment that platform design choices can cause harm through behavioral mechanisms, and not just through traditional market foreclosure. A platform that systematically redirects the attention of inattentive users toward its own services causes real harm. Ex ante rules that prohibit the design choices enabling that redirection are the right response.

## 6.2 Article 6(5) — Non-discriminatory ranking

Article 6(5) is the provision most directly relevant to the Google Shopping case. It requires gatekeepers not to treat their own services more favourably than those of third parties in rankings and search results. Google cannot use its results page to give Google Shopping an algorithmic advantage over competing comparison services. The obligation has three components: non-discrimination in ranking, transparency about ranking criteria, and fairness in the conditions applied to third parties.

The Commission has already acted on this provision. In March 2025, it issued a non-compliance decision finding that Google's search results page continued to favour its own services over competing ones, in violation of Article 6(5). This is significant. It confirms that the AT.39740 remedy did not solve the problem. Google's self-preferencing continued under a

new regulatory framework. It also shows that the DMA has real enforcement teeth: non-compliance decisions can lead to substantial fines.

From the behavioral perspective developed in this dissertation, Article 6(5) addresses the displacement mechanism discussed in Chapter 3. Systematically downranking competitors increases the effective search cost of reaching them. Prohibiting that downranking reduces that cost. But Article 6(5) says nothing about the visual format of results. A platform can comply with the ranking obligation while retaining a rich snippet format for its own services that is far more salient than what competitors can display. The ranking is neutral. The interface is not. This gap is examined in detail in Chapter 7.

## 6.3 Article 6(3) — Choice screens and defaults

Article 6(3) addresses default bias directly. It requires gatekeepers to allow users to easily change default settings that direct them toward the gatekeeper's own services. It also requires a choice screen to be shown to users at first use, allowing them to select their preferred search engine, browser, or virtual assistant.

This provision is a direct response to the evidence on default effects. The Commission identified Google's pre-installed default status on Android devices as a key source of its market power in AT.39740. The choice screen requirement tries to neutralise that advantage. By forcing an active choice at setup, Article 6(3) attempts to interrupt the passive behaviour that sustains default inertia.

The DMA's approach here was informed by economists who had studied default effects empirically. Decarolis, Li and Paternollo (forthcoming) documented a striking contrast: choice screens in the EEA reduced Google's market share by less than one percentage point, while mandatory default switches in Russia and Turkey reduced it by more than ten. Fletcher (2022) drew on this kind of evidence to argue that the DMA's approach to defaults was a genuine step forward but that its effectiveness would depend heavily on implementation. These contributions show that the choice screen provision was not designed blindly. It reflected real engagement with the behavioral economics literature.

Article 6(3) is, in this sense, the most explicitly behavioral provision in the DMA. It recognises that users do not actively choose their search engine in most cases. It tries to create a moment of deliberation where none existed before. Whether that moment is sufficient to change the behaviour of persistently inattentive users is a different question — one that Chapter 7 addresses.

## 6.4 Article 6(11) — Data access for competing search engines

Article 6(11) requires Google to share its search data, rankings, queries, clicks, and views, with competing search engines, on fair and non-discriminatory terms. The logic is clear. Google's dominance is partly self-reinforcing. More searches generate more data. More data improves the algorithm. A better algorithm attracts more users. And more users generate more searches. Competing engines cannot break into this cycle if they have no access to the data that drives it.

From a behavioral perspective, this provision addresses an indirect mechanism. Part of the H2 account of Google's dominance is that users choose Google because it genuinely performs better. And it performs better partly because it has more data. Forcing data sharing reduces this structural quality gap. It gives competing engines the tools to improve their results to a level where user choice can meaningfully operate. Without this, even perfect compliance with Article 6(5) might not restore effective competition. Users would rationally stick with Google because it is genuinely better.

Article 6(11) also has a methodological implication. Several of the experimental tests proposed in Chapter 4, particularly the post-click outcome study, would require access to data that Google currently holds exclusively. The data sharing obligation creates the legal basis for regulators to request that data as part of compliance monitoring.

## 6.5 The economists behind the DMA

The DMA did not emerge from a vacuum. Its provisions reflect a substantial body of academic work, and several of the economists behind that work appear directly in this dissertation's bibliography.

Fletcher, Crémer, Heidhues and co-authors (2024) document the role of economic analysis in shaping the DMA's design. Their assessment is that the DMA reflects genuine engagement with behavioral economics, particularly in its choice architecture provisions, but that the translation from research to regulatory text was imperfect in places. Fletcher (2022) had made a similar point earlier, arguing that the DMA's behavioral provisions were a good start but that their effectiveness would depend on whether they actually changed the behaviour of the users they were designed to protect.

Heidhues and Schnitzer (2021), writing at the time of the DMA negotiations, argued that the harm from Google's dominance was not purely structural. It was also cognitive. Structural obligations were necessary but not sufficient. Regulators needed to think about whether their remedies would reach the users whose behaviour they were trying to change.

These contributions are not just background context. They are evidence that the DMA was designed with behavioral harm in mind, at least partially. That makes the gaps in the DMA's behavioral coverage more interesting, not less. The question Chapter 7 asks is not whether the DMA tried to incorporate behavioral economics. It clearly did. The question is whether it went far enough.

The DMA is a real improvement over the ex post framework. It acts before harm is proven. It targets the design choices that enable self-preferencing. And it was shaped, at least in part, by economists who understood that digital markets involve behavioral harm, not just structural foreclosure. That is worth acknowledging. The critique that follows in Chapter 7 is not a rejection of the DMA. It is an argument that the DMA could do more, and a proposal for what that might look like.

# Chapter 7 — The Limits of the Digital Markets Act: What Behavioral Economics Reveals

## 7.1 A necessary but insufficient framework

The Digital Markets Act represents a genuine step forward. Chapter 6 showed that its core provisions, non-discriminatory ranking, choice screens, data sharing, were shaped by economists who understood that digital markets involve behavioral harm, not just structural foreclosure. That matters. It means the DMA is not starting from zero.

But a good start is not the same as a complete answer. Fletcher (2022), who contributed directly to the policy debate surrounding the DMA, puts it plainly: the DMA's behavioral provisions are a good start, but the Act arguably does not go far enough in considering the implications of behavioral economics. The following sections develop four specific limitations that substantiate that assessment. Each limitation corresponds to a behavioral mechanism documented in the earlier chapters of this dissertation, and each reveals a gap between what the DMA requires and what the behavioral evidence suggests would actually be needed.

The argument that follows is not a condemnation of the DMA. It is an invitation to improve it.

## 7.2 Limitation 1 — Article 6(5) regulates ranking but not visual salience

Article 6(5) requires Google to apply non-discriminatory ranking conditions. It prohibits the algorithmic self-preferencing documented in AT.39740, the systematic upranking of Google Shopping and downranking of competing services. That obligation is meaningful. The Commission's 2025 non-compliance decision shows it has real enforcement consequences.

But Article 6(5) says nothing about the visual format in which results are displayed. This is a significant gap. Google can comply fully with the ranking obligation while retaining a rich snippet format for its own services that is far more visually prominent than anything available to competitors. The ranking is neutral. The interface is not.

The distinction matters because of how visual salience operates. Chapter 3 reviewed the evidence from Bordalo, Gennaioli and Shleifer (2013) and Li and Camerer (2022). Salient attributes receive disproportionate attention automatically before deliberation begins. A result displayed with high-resolution images, gold star ratings, and prominent price information will capture more attention than a plain-text result at the same position on the page, regardless of the quality of the underlying content. Li and Camerer (2022) showed that this effect persists even when users are aware of it and try to correct for it. Deliberation reduces the salience effect but does not eliminate it.

Moreover, even for rational consumers, visual salience may cause harm through a pricing channel. As Kaye (2025) shows, concentrated demand enables above-competitive pricing. A visual format that systematically directs attention toward Google's services concentrates demand on those services, which in turn allows them to charge higher prices. The harm from visual salience is therefore not limited to fallible consumers.

The practical implication is direct. Google could theoretically respond to an Article 6(5) compliance requirement by adjusting its ranking algorithm while leaving the visual design of its Shopping units entirely unchanged. The click advantage that the Commission documented in AT.39740, approximately 42% of clicks going to Google Shopping, is driven by both the ranking position and the visual format.

There is a simple test for this claim. It is the visual ablation experiment proposed in Chapter 4: show Google Shopping results in a plain-text format identical to organic results, holding position and content constant. If the click advantage disappears, it was driven by visual salience, not by position alone. If it persists, it was driven by content quality. Article 6(5) does not require this test to be conducted. It does not require Google to demonstrate that its rich snippet format is informative rather than merely salient. And it does not prohibit Google from maintaining a visual design that systematically outperforms the formats available to competing services.

Closing this gap would not require a fundamental rewrite of Article 6(5). It would require an additional obligation: that the visual format of results displayed by the gatekeeper for its own services must not be more salient than the format available to competing services for equivalent queries. This is the logical extension of the non-discrimination principle that Article 6(5) already applies to ranking, applied to visual design rather than algorithmic position.

## 7.3 Limitation 2 — Choice screens do not reach persistently inattentive users

Article 6(3) requires Google to present users with a choice screen at first use, from which they can select their preferred search engine. The idea is sound. If users never actively consider which search engine they use, a choice screen creates a moment of deliberation that would not otherwise occur. It gives competing services a fair chance to be selected.

The problem is that the evidence does not support the assumption that most users will engage with that moment of deliberation. The most direct test comes from Allcott et al. (2025). Their field experiment included a condition in which users were presented with an active choice between Google and Bing before their search session began. The structural switching cost was zero: one click was all it took to select Bing. The result was an increase in Bing's market share of 1.1 percentage points.

This is a very small number. It implies that the vast majority of users, when given an explicit and frictionless choice, still select Google. The interpretation depends on which mechanism is driving that outcome. Under H2, it reflects genuine quality preference: users actively evaluate the options, conclude that Google is better, and choose it deliberately. Under H1, it reflects persistent inattention: users do not engage with the choice at all, select the default option, whichever it happens to be, and move on without deliberation.

The evidence from Decarolis et al. (forthcoming) helps distinguish these two interpretations. When the default was switched to Bing or a local competitor without any active choice requirement, market share effects were ten times larger than in the EEA choice screen intervention. If the small effect of the EEA choice screen reflected genuine quality preferences, the default switch should have produced a similarly small effect: users who

genuinely prefer Google would simply switch back. The fact that they did not suggests that inertia, not quality, is the dominant mechanism. And inertia is not cured by a choice screen.

Allcott et al. (2025) estimate that approximately one third of users are persistently inattentive. They never actively consider their choice of search engine, regardless of how easy switching has been made. For these users, a choice screen is not ineffective because it is poorly designed. It is ineffective because it requires a response that persistently inattentive users, by definition, do not provide. They see the screen. They click on the first or most familiar option. They move on.

Article 6(3) was designed for the two thirds of users who do deliberate. It does very little for the one third who do not. And that one third is not a small rounding error. It represents a large population of users whose behaviour sustains Google's market position through inertia rather than choice, and for whom the DMA's most explicitly behavioral provision provides essentially no protection.

What would actually reach these users? The evidence from Decarolis et al. (forthcoming) points toward a different type of intervention: not a choice screen that requires active engagement, but a mandatory periodic reassignment of defaults that forces exposure to alternatives regardless of user action. This is more intrusive than a choice screen. But the evidence suggests it is the only mechanism that actually changes the behaviour of persistently inattentive users. The DMA does not require it, and there is no provision in the current framework that would mandate anything like it.

## 7.4 Limitation 3 — The DMA is designed for rational consumers

This is the deepest limitation of the three discussed so far. It is not a gap in a specific provision. It is a gap in the underlying framework of the entire Act.

The DMA's obligations are written in the language of fairness and non-discrimination. Gatekeepers must apply fair and transparent ranking conditions. They must allow users to freely change their defaults. They must not treat their own services more favourably than those of third parties. These obligations make sense under one assumption: that users, once presented with equal conditions, will evaluate the available options and select the one that best serves their interests.

This is the rational consumer assumption. It is the same assumption that underlies standard competition law. And it is precisely the assumption that the behavioral economics literature has spent thirty years documenting is wrong, not for all users, and not in all circumstances, but systematically and predictably for a significant share of the population in routine, low-attention decisions.

Heidhues, Köster and Kőszegi (2023) make this point formally. Their model distinguishes between steering a rational consumer and steering a fallible consumer. When a rational consumer is steered toward a product, she evaluates it. If it is not what she wants, she rejects it. The welfare loss from steering a rational consumer is bounded, she retains the ability to correct for the steering once she engages with the choice. When a fallible consumer is steered, she does not evaluate. She adopts the steered product without comparison. The welfare loss is not bounded. As shown in Chapter 4, expressions (4.3) and (4.4) capture this formally: the

welfare loss for rational consumers under steering is small and bounded, while the loss for fallible consumers may be arbitrarily large.

The DMA's obligations do not reflect this distinction. Consider Article 6(5). The non-discrimination requirement assumes that competing services, once given equal ranking conditions, will attract users on the basis of quality. But for persistently inattentive users, equal ranking conditions are not sufficient. These users do not evaluate quality. They follow the most salient result, regardless of whether it is ranked fairly. Equal ranking without equal salience does not produce equal outcomes for inattentive users.

Consider Article 6(3). The choice screen assumes that users, once given an explicit choice, will evaluate it and select the best option. But for persistently inattentive users, a choice screen is just another interface element to click through. It does not trigger deliberation. It does not produce an informed choice. It produces a default selection dressed up as an active decision.

The pattern is consistent across provisions. The DMA consistently presupposes users who engage, evaluate, and choose deliberately. The behavioral evidence consistently shows that a significant share of users do not. The gap between these two pictures is not a detail. It is a structural mismatch between the model of the user embedded in the DMA and the actual behaviour of the users it is designed to protect.

Heidhues and Schnitzer (2021) made this point explicitly during the DMA negotiations. The harm from Google's dominance, they argued, was not purely structural, it was cognitive. Structural obligations were necessary but not sufficient. The regulators needed to think harder about whether their remedies would actually reach the users whose behaviour they were trying to change. That warning was not fully heeded in the final text of the DMA.

Addressing this limitation would require a fundamental shift in how the DMA conceptualises its target population. Rather than designing obligations for the median deliberating user, the DMA would need to design obligations for the full distribution of users, including the persistently inattentive third who never deliberate. That means moving from prohibitions on discriminatory design to positive requirements for interface design that does not exploit inattention. It means requiring gatekeepers to demonstrate that their compliance measures reach inattentive users, not just attentive ones. And it means accepting that some users need protection that goes beyond the removal of barriers to choice. They need an environment that does not exploit the fact that they are not choosing at all.

## 7.5 Limitation 4 — The effectiveness gap

The fourth limitation is, in some ways, the most practical. It concerns not what the DMA requires but how compliance is verified.

The DMA places significant emphasis on effectiveness. Gatekeepers must not just adopt compliant practices. They must ensure those practices actually work. This sounds like a behavioral commitment. In reality, the DMA delegates the responsibility for demonstrating effectiveness almost entirely to the gatekeepers themselves.

Under the current framework, Google designs its own compliance measures, implements them, and reports on them to the Commission. The Commission can investigate non-

compliance, as it did when it opened formal proceedings against Google in March 2024 and issued preliminary findings of non-compliance in March 2025. But the investigation is reactive: it starts from a complaint or from the Commission's own monitoring, and it requires the Commission to build a case that Google's measures are insufficient. There is no proactive requirement for Google to conduct independent behavioral testing of its compliance measures before they are implemented.

Fletcher et al. (2024) identify this as a central weakness. Where remedies involve measures designed to change consumer choices, their effectiveness will be affected by consumers' cognitive limitations and behavioral biases. The design of remedies will likely require empirical testing of various alternative approaches to demonstrate their effectiveness. However, the regulatory authorities may have neither the power to require the firm to conduct the necessary tests, nor economists with the required skills.

This is a striking admission from economists who contributed to the DMA's design. They are saying, in effect, that the DMA requires effectiveness but lacks the tools to verify it in the specific domain where verification is most necessary: the behavioral domain.

The practical consequence is visible in the choice screen evidence. The EEA choice screen intervention produced an effect of less than one percentage point. This is consistent with the behavioral prediction that active choice does not reach persistently inattentive users. But under the current DMA framework, a choice screen that produces a one percentage point effect could still be considered compliant, because the DMA does not specify what a behaviourally effective choice screen looks like or how its effectiveness should be measured. Google can present a choice screen, report that it was displayed to all users, and satisfy its obligation under Article 6(3) without any evidence that the screen changed the behaviour of the users it was supposed to reach.

Closing this gap would require three things. First, a definition of behavioral effectiveness: compliance measures must be shown to change the behaviour of the target population, including inattentive users, not just the behaviour of users who were already deliberating. Second, a requirement for independent testing: compliance measures in behavioral domains must be evaluated through independent controlled experiments, not just through metrics reported by the gatekeeper. Third, iterative review: if a compliance measure fails to demonstrate behavioral effectiveness, the gatekeeper must redesign it and test again, rather than treating initial compliance as a permanent discharge of the obligation.

None of these requirements exists in the current DMA. Their absence means that the most behaviorally informed regulation the EU has ever produced for digital markets may, in practice, change the behaviour of the users who needed no help to begin with, the attentive two thirds who were already capable of deliberate choice, while leaving the persistently inattentive third exactly where they were before the DMA came into force.

The four limitations documented in this chapter share a common root. The DMA was written for a world in which removing structural barriers to competition is sufficient to restore competitive outcomes. The behavioral evidence shows that this world does not exist, at least not for all users, and not in all circumstances. A third of users will not respond to choice screens. Inattentive users will not correct for visual salience once ranking conditions are equalised. Fallible consumers cannot protect themselves through deliberation that they never undertake. Effectiveness cannot be assumed from structural compliance alone.

These are not arguments against the DMA. They are arguments for taking the DMA's own stated commitment to effectiveness seriously, and for recognising that effectiveness, in a world of behaviorally heterogeneous users, requires more than the current text provides.

## Conclusion

This dissertation set out to answer one question. Does Google exploit behavioral biases to practice self-preferencing, or does it inform rational consumers navigating a complex information environment? The answer, after reviewing the theoretical and empirical literature, is not a simple yes or no. But it points clearly in one direction.

### *What the evidence shows*

The existing evidence is more consistent with H1 than H2. Three findings stand out.

First, approximately one third of Google's users are persistently inattentive. They never actively consider their choice of search engine. For these users, the default determines the outcome, and no amount of information about competing alternatives changes that. This finding, from Allcott et al. (2025), is the most direct empirical evidence available that a significant share of Google's market power operates through behavioral mechanisms rather than genuine quality superiority.

Second, the dominant source of user inertia is not structural switching costs. Allcott et al. (2025) show that active choice, which eliminates switching costs entirely, increases Bing's market share by only 1.1 percentage points. If switching costs were the main barrier, this number should be much larger. The small effect implies that the barrier is not technical. It is attentional.

Third, the visual design of Google's search results page, rich snippets, star ratings and high-resolution images, creates conditions of high visual salience that capture attention before rational evaluation can begin. Bordalo, Gennaioli and Shleifer (2013) provide the theoretical mechanism, and Li and Camerer (2022) confirm experimentally that this effect persists even among users who are aware of it and try to correct for it.

Together, these findings support the behavioral exploitation account of self-preferencing. But they do not constitute a decisive test. The central gap in the literature is the absence of an experiment that holds product quality constant while varying the visual format of Google Shopping results. Without this, it is impossible to determine whether Google Shopping's click advantage reflects the informational value of its visual design or simply the attentional capture that salience generates independently of content. This dissertation has proposed four experimental designs that would close that gap. Conducting them and making the results publicly available is the most important next step in this research agenda.

The evidence that is available, however, is sufficient to raise serious concerns. Google's market position is not fully explained by quality. Persistent inattention plays a large and independent role. And the behavioral mechanisms through which self-preferencing operates, attention capture, default inertia and visual salience, are not addressed by the structural remedies that competition law has traditionally relied on.

### *What the DMA gets right and where it falls short*

The Digital Markets Act is a genuine step forward. It acts before harm is proven. It imposes obligations on Google's ranking practices, its default settings, and its data sharing, all of which correspond to real mechanisms of harm identified in the behavioral literature. And it was shaped, at least in part, by economists who understood that digital markets involve behavioral harm, not just structural foreclosure.

But the DMA falls short in four specific ways, documented in Chapter 7.

It regulates ranking but not visual salience. A gatekeeper can comply fully with Article 6(5) while maintaining a rich snippet format that is systematically more salient than what competing services can display. The non-discrimination principle that the DMA applies to algorithmic ranking does not extend to visual design.

It relies on choice screens that do not reach persistently inattentive users. Article 6(3) creates a moment of deliberation where none existed before. But deliberation requires users who are willing to deliberate. For the one third of users who never actively consider their options, a choice screen is just another interface element to click through without thought.

It is designed for rational consumers. All of its core obligations, non-discriminatory ranking, easy default modification and transparent criteria, make sense if users evaluate options carefully once the structural conditions are equalised. The behavioral evidence shows that a significant share of users do not. The DMA has no specific answer for persistently inattentive users who cannot protect themselves through deliberation they never undertake.

And it delegates effectiveness to the gatekeeper. Fletcher et al. (2024) acknowledge the problem directly: compliance measures in behavioral domains need empirical testing to demonstrate they actually work, but the current framework places the responsibility for that testing on the gatekeeper itself. Google decides whether its compliance measures are behaviourally effective. There is no requirement for independent verification.

### *Two concrete recommendations*

This dissertation closes with two specific recommendations for improving the DMA's behavioral coverage.

The first is to extend the non-discrimination obligation of Article 6(5) to visual format, not just algorithmic ranking. Concretely, this means requiring that the visual format of results displayed by the gatekeeper for its own services must not be more salient than the format made available to competing services for equivalent queries. This is the logical extension of a principle already present in the DMA. Applying it to visual design would close the gap between the ranking neutrality the DMA requires and the interface neutrality that behavioral evidence suggests is necessary.

The second is to impose mandatory independent behavioral testing of compliance measures. Before any compliance measure in a behavioral domain is accepted as sufficient, it should be required to demonstrate, through an independent randomised controlled experiment, that it changes the behaviour of the target population, including persistently inattentive users, not just attentive ones. If the measure fails to demonstrate effectiveness, the gatekeeper must

redesign and retest. This shifts the burden of proof from the regulator to the gatekeeper, and it ensures that compliance means something more than the adoption of a formally correct but behaviourally ineffective remedy.

These two recommendations are modest in scope but significant in principle. They do not require a rewrite of the DMA. They require that the DMA take its own commitment to effectiveness seriously, and that it recognise, explicitly and formally, that effectiveness in a world of behaviorally heterogeneous users is not guaranteed by structural remedies alone.

### *Limitations of this dissertation*

This dissertation has three main limitations worth acknowledging.

The empirical evidence it draws on comes primarily from the United States and from platform contexts, Expedia, eBay and desktop search, that are not identical to the Google Shopping context in the EU. The behavioral mechanisms are likely similar across contexts, but the magnitudes may differ. Estimates of the share of persistently inattentive users, in particular, may not transfer directly from a US desktop search experiment to EU mobile users.

The analysis is qualitative. This dissertation does not produce original data or run original statistical tests. Its contribution is analytical: it connects theoretical frameworks, empirical findings from the existing literature, and a regulatory analysis of the DMA into a coherent argument. That argument is only as strong as the evidence it draws on, and some of that evidence, particularly the experimental literature on position effects in Google Shopping specifically, is limited.

Finally, the DMA is a young regulation. Its implementation is ongoing, and its effects on market outcomes are not yet fully observable. The compliance decisions issued in 2025 suggest real enforcement activity. Whether that activity will produce meaningful behavioral change in the population of persistently inattentive users is a question that only future research can answer.

### *A final thought*

The Google Shopping case proved that self-preferencing causes competitive harm. It took seven years to do so, and the conduct it addressed continued after the decision. The Digital Markets Act was designed to move faster and act earlier. It is better than what came before.

But speed and early action are not enough if the underlying model of the user is wrong. A regulation designed for deliberating, rational consumers will protect the users who needed the least help, those who were already capable of navigating the market on their own terms. It will do very little for the users who need it most: the third who never deliberate, the users whose attention is captured before evaluation begins, the users for whom the default is not a starting point but an endpoint.

Getting digital regulation right means designing it for the users who actually exist, not the users that standard economic models assume. That is what this dissertation has argued. And it is, ultimately, what the behavioral economics literature that Heidhues, Kőszegi, Fletcher, and others have built over the past two decades has been trying to tell regulators all along.

## Bibliography


Akman, P. (2022). The theory of abuse in Google Search. *Journal of Antitrust Enforcement*, *10*(1), 1–34.

Allcott, H., Castillo, J. C., Gentzkow, M., Musolff, L., & Salz, T. (2025). Sources of market power in web search: Evidence from a field experiment (NBER Working Paper No. 33410). National Bureau of Economic Research.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, *121*(5), 803–843.

Crémer, J., Dinielli, D., Heidhues, P., Kimmelman, G., Monti, G., Podszun, R., Schnitzer, M., Scott Morton, F., & de Streel, A. (2023). Enforcing the Digital Markets Act: Institutional choices, compliance, and antitrust. *Journal of Antitrust Enforcement*, *11*(3), 315–349.

Crémer, J., Kimmelman, G., Monti, G., Podszun, R., Schnitzer, M., & Scott Morton, F. (2024, February 27). What we learn about the behavioral economics of defaults from the Google Search monopolization case. *ProMarket*. https://www.promarket.org/2024/02/27/what-we-learn-about-the-behavioral-economics-of-defaults-from-the-google-search-monopolization-case/

De los Santos, B., Hortaçsu, A., & Wildenbeest, M. R. (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, *102*(6), 2955–2980.

Decarolis, F., Li, M., & Paternollo, F. (Forthcoming). Competition and defaults in online search. *American Economic Journal: Microeconomics*. (CEPR Discussion Paper No. 17779)

Dinerstein, M., Einav, L., Levin, J., & Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, *108*(7), 1820–1859.

European Commission. (2017). *Case AT.39740 — Google Search (Shopping)*. Official Journal of the European Union.

European Commission. (2025, March 19). *Commission sends preliminary findings to Alphabet under the Digital Markets Act*. https://digital-markets-act.ec.europa.eu/commission-sends-preliminary-findings-alphabet-under-digital-markets-act-2025-03-19_en

Fletcher, A. (2022, October 31). Behavioral insights in the DMA: A good start, but how will the story end? *Competition Policy International*. https://www.competitionpolicyinternational.com/behavioral-insights-in-the-dma-a-good-start-but-how-will-the-story-end/

Fletcher, A., Crémer, J., Heidhues, P., Kimmelman, G., Monti, G., Podszun, R., Schnitzer, M., Scott Morton, F., & de Streel, A. (2024). The effective use of economics in the EU Digital Markets Act. *Journal of Competition Law and Economics*, *20*(1–2), 1–19.

Haan, M. A., Moraga-González, J. L., & Petrikaitė, V. (2018). A model of directed consumer search. *RAND Journal of Economics*, *49*(2), 377–401.
Heidhues, P., & Kőszegi, B. (2018). Behavioral industrial organization. In B. D. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of Behavioral Economics* (Vol. 1, pp. 517–612). Elsevier.

Heidhues, P., & Schnitzer, M. (2021, October 5). Nur Google profitiert von Googles Dominanz [Op-ed]. *Handelsblatt*. https://www.handelsblatt.com/meinung/gastbeitraege/gastkommentar-nur-google-profitiert-von-googles-dominanz/27672774.html

Heidhues, P., Köster, M., & Kőszegi, B. (2023). Steering fallible consumers. *The Economic Journal*, *133*(652), 1430–1467.

Heidhues, P., Köster, M., & Kőszegi, B. (2024). *A theory of digital ecosystems* [Working paper].

Kaye, A. P. (2025). *The personalization paradox: Welfare effects of personalized recommendations in two-sided digital markets* [Working paper]. MIT Sloan School of Management.

Kittaka, Y., Sato, S., & Zennyo, Y. (2023). Self-preferencing by platforms: A literature review. *International Journal of Industrial Organization*, *88*, 102944.

Kohavi, R., Tang, D., & Xu, Y. (2020). Trustworthy online controlled experiments: A practical guide to A/B testing. Cambridge University Press.

Li, X., & Camerer, C. F. (2022). Predictable effects of visual salience in experimental decisions and games. *Quarterly Journal of Economics*, *137*(3), 1849–1900.

Liu, L., & Dukes, A. (2016). Consumer search with limited product evaluation. *Journal of Economics & Management Strategy*, *25*(1), 5–30.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, *69*(3), 213–225.

Ursu, R. M. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, *37*(4), 530–552.

Vásquez Duque, O. (2022). *The potential anticompetitive stickiness of default applications: Addressing consumer inertia with randomization* [Working paper]. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4077132