



**UNIVERSITY
OF TURKU**

Turku School of
Economics

Privacy or Performance? The GDPR's Impact on Online Advertising

Information Systems Science,
Department of Management and Entrepreneurship
Master's thesis

Author:
Andreas Weinberger

Supervisors:
Professor Jukka Heikkilä
D.Sc Eija Koskivaara

13.02.2026
Turku

Student's statement regarding the use of Artificial Intelligence (AI) for preparing and/or writing this thesis:

☐ **I have not used any AI-based tools.**

☒ **I have used AI-based tools.** Their use is documented in the Appendix. The AI tools were used in a way that complies with academic integrity guidelines.

The originality of this thesis has been checked in accordance with the University of Turku quality assurance system using the Turnitin Originality Check service.

Master's thesis

Subject: Information Systems Science

Author(s): Andreas Weinberger

Title: Privacy or Performance? The GDPR's Impact on Online Advertising

Supervisor: D.Sc Eija Koskivaara

Number of pages: 48 pages (+ appendice 1 page)

Date: 13.02.2026

Abstract

This thesis reviews empirical evidence on the effects of the General Data Protection Regulation (GDPR) on online advertising markets. Focusing on advertising-funded business models, it documents how the regulation altered firms' data practices and examines how these changes affected user privacy, advertising effectiveness, and economic efficiency. The analysis further considers the implications of these effects for market structure and competition. The literature shows that the GDPR substantially reduced user observability by limiting tracking, identifier persistence, and cross-firm data sharing. These privacy gains arise mainly through compliance-driven retrenchment and user opt-out rather than through effective user control, and they are unevenly distributed across users. At the same time, reduced access to personal data weakened advertising effectiveness. Restrictions on targeting and measurement lowered advertiser participation and revenues for many publishers, though losses were partially mitigated by substitution toward contextual targeting. Finally, the GDPR reshaped competitive dynamics. Large platforms with strong first-party data adapted more successfully than smaller publishers and intermediaries, contributing to increased concentration and reduced entry in data-intensive markets. Overall, the findings highlight a trade-off between improved privacy protection and efficiency costs in the ad-supported digital economy.

Keywords: General Data Protection Regulation (GDPR), Online Advertising Markets, Privacy, Efficiency, Data Practices, Tracking, Advertising Effectiveness, Market Concentration, Consent

TABLE OF CONTENTS

1	Introduction	7
2	Background: The Online Advertising Ecosystem and the GDPR	9
	2.1 The Advertising Value Chain	9
	2.2 Identifiers, Cross-Firm Tracking, and Data Types	9
	2.3 Real-Time Bidding as a Data Infrastructure	11
	2.4 GDPR Principles Relevant to Advertising and Data Flows	11
3	Conceptual Framework: Privacy, Observability, Ad Performance, and Welfare .	14
	3.1 Privacy and Observability in Online Advertising	14
	3.2 Consent, Selection, and the Recording Wedge	15
	3.3 Advertising Performance and Measurement	15
	3.4 From Advertising Performance to Revenue and Welfare	16
4	GDPR and Changes in Data Practices	18
	4.1 Mechanism: Risk-Based Retrenchment	18
	4.2 Evidence on Changes in Data Volumes, Intensity, and Flows	19
	4.2.1 Backend Data Storage and Processing	19
	4.2.2 Client-Side Tracking and Identifier Availability	20
	4.2.3 Cross-Firm Data Sharing and Network Structure	20
5	Effects of GDPR on User Privacy	22
	5.1 Privacy as Reduced Observability	22
	5.2 Limits of Control and Consent	23
	5.2.1 Transparency Without Comprehension	23
	5.2.2 Consent Design and Behavioral Constraints	24
	5.2.3 Technical Limits to Enforcing Data Rights	24
	5.2.4 Implications for User Control	25
	5.3 Heterogeneous Privacy Effects and Externalities	25
	5.3.1 Privacy Sensitivity and Unequal Protection	25
	5.3.2 Coarse Consent, Sorting, and Privacy Externalities	26
	5.3.3 Data Deserts and Exclusion Trade-Offs	27
6	Effects of GDPR on Efficiency and Firm Performance	28
	6.1 Direct Effects on Advertising Effectiveness	28

6.1.1	Baseline Value of Data-Driven Targeting	28
6.1.2	Evidence from Publisher-Side Data	29
6.1.3	Advertising as a Traffic-Generation Mechanism	30
6.1.4	Evidence from Intermediary Data	30
6.1.5	Reconciling Evidence Across Institutional Settings	31
6.1.6	Contextual Targeting and Selection Effects	31
6.2	Effects on Firm Performance and Market Structure	32
6.2.1	From Advertising Performance to Firm-Level Revenue and Profitability	32
6.2.2	Differential Impacts and Scale Effects	33
6.2.3	Content Supply and Output Responses	34
6.2.4	Market Structure and Concentration	35
6.2.5	Entry, Innovation, and Dynamic Efficiency	36
6.3	Actor-Specific Adjustments and Implications	37
6.3.1	Publishers: Compliance Layering and Risk Reallocation	37
6.3.2	Advertisers: Reallocation Toward Reliable Attribution	38
6.3.3	Intermediaries: Consolidation and Gatekeeping	39
7	Privacy and Performance under the GDPR: A Synthesis	40
7.1	Intended vs. Realized Privacy: Data Minimization without Effective User Control	40
7.2	Efficiency Costs of Privacy Regulation	42
7.3	Redistribution across the Advertising Value Chain	44
8	Limitations and Future Research	47
9	Conclusion	48
	References	49
	Appendices	55
	Appendix 1 Explanation of the use of AI	55

1 Introduction

Online advertising is a central component of the digital economy, financing content and services that are often offered to users at zero monetary price (Chen, 2022; Congiu, Sabatino, & Sapi, 2022). Its effectiveness relies on the collection and linkage of user-level data across time, websites, and firms, enabling behavioral targeting, performance measurement, and learning (Athey & Gans, 2010; Miller, Lukic, & Skiera, 2025; W3C Tracking Protection Working Group, 2019). While these practices improve the precision of matching ads to users, they have raised concerns about transparency, autonomy, and control over personal data (Acquisti, Taylor, & Wagman, 2016; Eggers, Beke, Verhoef, & Wieringa, 2023). Against this background of growing concerns about data-intensive digital business models, the European Union implemented the General Data Protection Regulation (GDPR), adopted in 2016 and enforced from May 2018 (European Parliament and Council of the European Union, 2016). The GDPR introduced stricter requirements for the collection, processing, and sharing of personal data, including explicit and informed user consent for many forms of behavioral advertising and cross-site tracking (Ke & Sudhir, 2023; Nokhbeh Zaeem & Barber, 2020). For online advertising, one of the most data-intensive sectors of the digital economy (Congiu et al., 2022; Peukert, Bechtold, Batikas, & Kretschmer, 2022), this represents a direct intervention into the data flows underlying ad targeting, attribution, and real-time bidding.

A growing empirical literature shows that these constraints have economically meaningful consequences. Limits on persistent identifiers and cross-firm data sharing reduce advertisers' ability to target users and measure campaign effectiveness, weakening ad-funded revenue models particularly for publishers and intermediaries that rely heavily on programmatic advertising (Johnson, Shriver, & Du, 2020; P. Wang, Jiang, & Yang, 2024). At the same time, firms with access to extensive first-party data or operating within vertically integrated ecosystems appear better positioned to adapt, raising concerns about increased concentration and reduced competition (Aridor, Che, & Salz, 2023; Janßen, Kesler, Kummer, & Waldfogel, 2022; Johnson, Shriver, & Goldberg, 2023; Peukert et al., 2022). These developments point to a central trade-off between enhanced user privacy and economic performance in online advertising markets (Choe, Matsushima, & Shekhar, 2025).

This thesis focuses on the online advertising market and provides a systematic synthesis of the empirical literature on the GDPR's effects. It examines how the regulation altered data practices,

user privacy outcomes, advertising performance, and market structure. The thesis addresses three research questions:

- To what extent has the GDPR improved user privacy in online advertising?
- How has the GDPR affected the performance of online advertising and ad-funded revenue models?
- What are the broader implications for market structure and the distribution of revenues between large and small firms?

The remainder of this thesis is structured as follows. Chapter 2 introduces the online advertising value chain and relevant GDPR principles, and Chapter 3 develops the conceptual framework. Chapter 4 examines how data practices changed following the GDPR, while Chapters 5 and 6 analyze how these changes affected user privacy and economic efficiency. Chapter 7 synthesizes and interprets these findings. Chapter 8 discusses limitations and directions for future research, and Chapter 9 concludes.

2 Background: The Online Advertising Ecosystem and the GDPR

This chapter provides the institutional, technical, and legal background necessary to analyze the effects of the GDPR on online advertising markets. It first provides a high-level overview of the advertising value chain and its key components, and then examines in more detail how user data are generated, linked, and transmitted through identifiers and real-time bidding. The chapter concludes by outlining the GDPR principles that constrain these data flows.

2.1 The Advertising Value Chain

Online advertising is a socio-technical system that sells user attention at scale and in real time and relies on interactions between users, publishers, advertisers, and intermediaries to generate, transmit, and monetize user data. These intermediaries include supply-side platforms (SSPs), demand-side platforms (DSPs), ad exchanges, and data management platforms (DMPs), which facilitate matching, bidding, and data processing (Urban, Tatang, Degeling, Holz, & Pohlmann, 2020). In addition, specialized tracker providers supply embedded software that bundles functional services (e.g., analytics, social sharing) with data collection capabilities (Miller et al., 2025). Following the introduction of privacy regulation, consent management platforms (CMPs) emerged as an additional infrastructure layer, collecting user consent and transmitting standardized consent signals to downstream advertising and measurement systems, thereby mediating access to user data across publishers and third-party vendors (Hils, Woods, & Böhme, 2020; Johnson et al., 2023).

The functioning of this ecosystem relies on user data generated through interaction with digital services. These data enable targeting, measurement, and learning, and circulate across publishers, intermediaries, and tracker providers. The technical mechanisms underlying user identification and cross-firm data sharing are examined in Section 2.2.

Most display advertising transactions are coordinated through Real-Time Bidding (RTB), which serves as the core market infrastructure linking data flows to real-time allocation and pricing decisions (Veale & Zuiderveen Borgesius, 2022; J. Wang, Zhang, & Yuan, 2017). RTB is discussed in detail in Section 2.3.

2.2 Identifiers, Cross-Firm Tracking, and Data Types

This section explains the technical mechanisms through which users are identified and tracked in online advertising and clarifies why these mechanisms are central to cross-firm data sharing and regulatory intervention. Tracking technologies are the primary means through which user-level data

are generated in online advertising. These technologies are embedded in websites and applications to support functions such as analytics, measurement, personalization, and advertising, though not all tracking is advertising-relevant. While some tracking supports essential functionality or security, other forms enable behavioral advertising by recording and linking user activity across sessions and contexts (Miller et al., 2025; W3C Tracking Protection Working Group, 2019).

Because the web's underlying HTTP protocol is stateless, online advertising relies on persistent identifiers to recognize users across visits and contexts. Technologies such as cookies and mobile advertising identifiers enable firms to link user interactions over time (Rafieian & Yoganarasimhan, 2021; Veale & Zuiderveen Borgesius, 2022; J. Wang et al., 2017).

Beyond recognition within a single website or application, advertising systems depend on linking identifiers across organizational boundaries. While browsers enforce a same-origin policy that restricts data access to the domain that set it, advertising technology firms historically enabled cross-firm recognition through identifier synchronization (also known as cookie syncing), in which user identifiers are exchanged and mapped across firms (Acar et al., 2014; Urban et al., 2020). This process allows user-level data to circulate across networks of advertising and analytics providers, supporting targeting, measurement, and attribution across websites and applications.

A central distinction in these data flows is between first-party and third-party data. First-party data are collected directly by a publisher or platform from users interacting with its own service or voluntarily providing information, such as registration details or account attributes, within a direct user–firm relationship (Godinho de Matos & Adjerid, 2022; Peukert et al., 2022). Third-party data, by contrast, are collected by entities without a direct relationship to the user, typically through embedded tracking technologies or data-sharing arrangements across organizational boundaries (Peukert et al., 2022).

This distinction is legally and operationally important. First-party data are generally easier to govern, justify, and explain to users, as they are collected within a clear service context and often rely on a single data controller. Third-party data, in contrast, depend on complex consent chains and involve multiple data controllers, increasing compliance risk and coordination costs (Veale & Zuiderveen Borgesius, 2022). Because online advertising historically relied heavily on third-party data for cross-site tracking and behavioral targeting, this distinction plays a key role in shaping how privacy regulation redistributes data access across firms (Johnson et al., 2023).

2.3 Real-Time Bidding as a Data Infrastructure

Real-time bidding (RTB) is the core transaction mechanism of modern display advertising, in which impressions are auctioned automatically and contextual and user-level data are translated into allocation and pricing decisions in real time (Urban et al., 2020; Veale & Zuiderveen Borgesius, 2022; Yuan, Wang, & Zhao, 2013).

Most RTB environments use generalized second-price auctions, in which advertisers bid based on their expected value from showing an ad to a given user or context, and the winning bidder pays a price determined by the next-highest bid (Edelman, Ostrovsky, & Schwarz, 2007; Varian, 2007). Observed bids and prices therefore reflect advertisers' assessments of user value, targeting precision, and attribution accuracy (J. Wang et al., 2017).

Because bid requests embed contextual information and, when permitted, signals derived from persistent identifiers, RTB integrates data flows directly into allocation and pricing decisions. As a result, restrictions on data collection, identifier persistence, or cross-firm data sharing can directly affect auction outcomes, advertiser participation, and publisher revenue (Goldberg, Johnson, & Shriver, 2024; Johnson et al., 2020; P. Wang et al., 2024).

2.4 GDPR Principles Relevant to Advertising and Data Flows

This section outlines the GDPR principles most relevant for online advertising and explains how they are intended to reshape firms' data practices and strengthen users' control over personal data. Prior to the GDPR, personal data processing in the EU was governed by the Data Protection Directive (95/46/EC) and the ePrivacy Directive (2002/58/EC) (European Parliament and Council of the European Union, 1995, 2002). Although the latter required user consent for storing information on user devices, enforcement was fragmented and implementation varied widely across member states, leading firms to rely in practice on implied or opt-out consent (Congiu et al., 2022; Degeling et al., 2019; Veale & Zuiderveen Borgesius, 2022). The GDPR replaced the Data Protection Directive to harmonize data protection rules and significantly tightened the definition of valid consent, explicitly precluding silence or inactivity as a lawful basis for processing personal data (Veale & Zuiderveen Borgesius, 2022).

Beyond harmonizing data protection rules, a central normative objective of the GDPR is to strengthen individuals' control over their personal data. The regulation explicitly emphasizes that natural persons should have control over information relating to them (GDPR Recital 7) and seeks

to operationalize this objective through stricter consent requirements, enhanced transparency obligations, and expanded data subject rights. In this sense, the GDPR is not only a constraint on firms' data practices but also an attempt to reallocate decision rights over personal data toward users, particularly in contexts characterized by opaque data flows and asymmetric information such as online advertising.

Against this background, several GDPR principles are particularly relevant for online advertising and data flows. First, the principles of data minimization and purpose limitation require that personal data be adequate, relevant, and limited to what is necessary for specified purposes (GDPR Art. 5; Nokhbeh Zaeem & Barber, 2020). These principles challenge tracking-by-default models that collect broad user data for multiple downstream uses, including targeting, attribution, and profiling, and restrict the repurposing of data across advertising functions without renewed legal justification.

Second, the GDPR requires a lawful basis for processing personal data (GDPR Art. 6) and explicitly classifies online identifiers, such as cookie IDs and IP addresses, as personal data (GDPR Art. 4). In advertising contexts, regulators generally maintain that behavioral advertising and cross-site tracking require explicit, informed, and unambiguous user consent (GDPR Art. 4(11) & Art. 7; Ke & Sudhir, 2023; Veale & Zuiderveen Borgesius, 2022). While the GDPR also permits processing based on "legitimate interest" (GDPR Art. 6(1)(f)), legal scholarship emphasizes that this basis is subject to a strict three-part test: the processing must pursue a legitimate interest, it must be necessary to achieve that interest, and it must not be overridden by the fundamental rights and freedoms of the user (Veale & Zuiderveen Borgesius, 2022). In the context of behavioral advertising, scholars and regulators generally argue that the intrusiveness of tracking fails this balancing test, implying that consent is the only valid legal basis (Veale & Zuiderveen Borgesius, 2022).

Third, transparency obligations and data subject rights require firms to disclose which entities process personal data and for what purposes, and to enable rights such as access, erasure, and restriction of processing. These requirements are particularly demanding in advertising ecosystems characterized by multiple intermediaries and complex data-sharing relationships, where informing users about all downstream recipients and purposes is technically challenging (Nokhbeh Zaeem & Barber, 2020; Veale & Zuiderveen Borgesius, 2022).

Finally, the GDPR imposes substantial penalties for non-compliance, reaching up to 20 million euros or 4% of a firm's total global annual turnover, whichever is higher (GDPR Art. 83(5); Peukert

et al., 2022). In addition, European court rulings have expanded notions of joint controllership, implying that websites embedding third-party technologies may be held jointly liable for downstream data misuse. This shared liability increases the legal and reputational risks associated with complex tracking and cross-firm data-sharing arrangements and creates strong incentives for publishers to vet, limit, or simplify their use of third-party vendors (Peukert et al., 2022; Veale & Zuiderveen Borgesius, 2022).

3 Conceptual Framework: Privacy, Observability, Ad Performance, and Welfare

This chapter introduces the conceptual framework used to organize the empirical evidence reviewed in subsequent chapters. The framework focuses on three elements: privacy as observability, consent-based selection, and the measurement of advertising performance.

3.1 Privacy and Observability in Online Advertising

This section defines how privacy is conceptualized in the thesis and introduces observability as the key analytical construct. In the digital economy, privacy is commonly understood as control over the collection, reuse, and circulation of personal information rather than as a binary condition of secrecy (Acquisti et al., 2016; Westin, 1967). From this perspective, privacy concerns arise when information about individuals is monitored, linked, or repurposed across contexts beyond their reasonable expectations. A prominent normative account is Contextual Integrity, which defines privacy as the appropriateness of information flows relative to contextual purpose and social norms (Nissenbaum, 2010).

In online advertising, these normative concerns map naturally to the extent to which firms can observe user behavior. Accordingly, this thesis operationalizes privacy in terms of observability: the degree to which firms can monitor, link, and retain information about individual users across contexts and over time (Aridor et al., 2023; Goldberg et al., 2024). Higher observability corresponds to richer and more persistent user-level data, while lower observability reflects constraints on tracking, linkage, and data retention.

This perspective clarifies why privacy regulation reshapes data flows rather than eliminating data use altogether. Users are more willing to accept first-party data collection aligned with a service context than third-party tracking that repurposes data across contexts (Nissenbaum, 2010), implying asymmetric effects on first- and third-party data. In online advertising, observability is embedded in market infrastructure such as real-time bidding, which integrates contextual and user-level information directly into allocation and pricing decisions (Veale & Zuiderveen Borgesius, 2022). Privacy regulation therefore affects advertising outcomes both by constraining which user behavior can be recorded and by determining which users become observable through consent.

3.2 Consent, Selection, and the Recording Wedge

This section explains how consent-based privacy regimes translate heterogeneous privacy preferences into selection between observable and unobservable users, shaping what user behavior is observed in advertising markets. Consumers differ substantially in how they value privacy. Economic research distinguishes between intrinsic privacy preferences, where tracking is experienced as inherently undesirable, and instrumental preferences, which reflect concerns about downstream economic consequences such as price discrimination (Lin, 2022). These differences imply that consent decisions are systematically related to user characteristics that are economically relevant for advertising.

Under regimes such as the GDPR, users can often opt out of tracking while continuing to access digital services. As a result, users sort into observable and unobservable populations. Firms and researchers observe only the behavior of users who consent to tracking, while the behavior of non-consenting users remains unrecorded (Aridor et al., 2023; Johnson et al., 2020). Observed data therefore reflect a selected subset of users rather than a random sample of overall demand.

This sorting creates a divergence between actual economic activity and recorded outcomes. Goldberg et al. (2024) refer to this divergence as a recording wedge: observed data reflect only the behavior of consenting users, while real outcomes include the full population. Because consent decisions are correlated with user characteristics such as engagement or purchase propensity, this wedge introduces non-random selection that can distort the informational content of observed data (Aridor et al., 2023; Goldberg et al., 2024). Importantly, the recording wedge affects not only empirical inference but also firms' real-time decision-making, as targeting, learning, and attribution are all based on the behavior of the observable subset.

At the aggregate level, individual consent decisions may also generate information externalities. Data provided by some users can improve inference about others, implying that privacy choices affect the overall informational environment faced by firms (Choi, Jeon, & Kim, 2019). As a result, consent-based selection influences advertising effectiveness and welfare not only through who is observed, but also through how informative observed data are for predicting unobserved behavior.

3.3 Advertising Performance and Measurement

This section defines how advertising performance is measured in online advertising markets and explains why observability is central to interpreting these performance metrics. Advertising

performance is assessed using metrics that capture both the effectiveness of individual ad exposures and the scale of advertising activity. Common per-exposure measures include click-through rates (CTR), conversion rates, and attribution metrics that allocate credit for conversions across a user's exposure path (J. Wang et al., 2017; P. Wang et al., 2024). Measures of scale include total impressions, total clicks, and aggregate advertising revenue (Goldberg et al., 2024; P. Wang et al., 2024).

From the publisher perspective, performance is reflected in bid prices, advertiser participation, and revenue per impression or click, which together capture demand for access to a given audience (Urban et al., 2020; P. Wang et al., 2024).

All of these measures depend on observability. Persistent identifiers enable targeting, frequency capping, learning, and attribution by linking ad exposures to user behavior over time (Johnson et al., 2020; J. Wang et al., 2017). When observability is constrained, advertising performance metrics become less informative about targeting precision, learning dynamics, and attribution accuracy, potentially weakening the link between observed outcomes and underlying user behavior (Goldberg et al., 2024; Rafieian & Yoganarasimhan, 2021; Urban et al., 2020; P. Wang et al., 2024).

3.4 From Advertising Performance to Revenue and Welfare

This section links advertising performance to firm revenue and economic welfare by clarifying how changes in observability and measurement translate into allocative and dynamic efficiency outcomes. Advertising performance provides the central link between data availability, firm revenue, and welfare in online advertising markets. Richer user-level data improve advertisers' ability to target ads and measure their effectiveness, increasing expected returns from advertising activity (Rafieian & Yoganarasimhan, 2021; P. Wang et al., 2024). Improved observability enables more precise matching between ads and users and more accurate attribution of outcomes to advertising exposures, reducing measurement noise and uncertainty about campaign performance (Aridor et al., 2023). As a result, advertisers' willingness to pay for advertising impressions increases.

In auction-based advertising markets, higher expected returns are reflected in higher bids and higher prices for access to well-measured or identifiable audiences (Farahat & Bailey, 2012; Johnson et al., 2020). Improvements in targeting and measurement therefore translate into higher aggregate advertising revenue, while constraints on data availability weaken targeting precision and attribution, lowering advertisers' expected returns and demand for advertising inventory.

Advertising revenue, in turn, finances the provision of digital content and services that are often offered to users at zero monetary price (Anderson & Gabszewicz, 2006). Beyond covering variable operating costs, advertising income supports investment in product development, experimentation, and market entry, linking advertising performance to innovation and dynamic efficiency (Chen, 2022; Janßen et al., 2022; Jia, Jin, & Wagman, 2021). Changes in advertising revenue therefore affect welfare not only through short-run allocative efficiency, but also through longer-run incentives to invest and supply new products or content.

From a welfare perspective, reduced targeting precision can lower allocative efficiency by weakening the match between ads and consumers (Goldfarb & Tucker, 2011; Rafieian & Yoganarasimhan, 2021), while reduced advertising revenue can constrain the financing of content and innovation (Janßen et al., 2022; Jia et al., 2021). At the same time, theory highlights that the welfare effects of privacy regulation depend on firms' revenue models: opt-in regimes can increase total welfare when firms rely primarily on usage-based rather than data-based revenues, as expanded participation by privacy-sensitive users can outweigh losses from reduced data monetization (Choe et al., 2025). Advertising markets nevertheless remain central to subsidizing access to digital services, implying that changes in advertising performance may affect consumer surplus even when monetary prices remain zero (Anderson & Gabszewicz, 2006; Dubé et al., 2025).

By shaping data availability, observability, and measurement accuracy, privacy regulation can therefore affect economic welfare indirectly through its impact on advertising performance and revenue (Acquisti et al., 2016). The subsequent chapters use empirical evidence to assess how these mechanisms operate in practice and to evaluate the resulting trade-offs between privacy protection, efficiency, and welfare (Jones & Tonetti, 2020).

4 GDPR and Changes in Data Practices

This chapter examines how firms adjusted data practices in online advertising following the enforcement of the GDPR. It documents changes in data collection, processing, retention, and cross-firm sharing, and shows that enforcement triggered an initial contraction in data flows followed by partial adaptation and reconfiguration rather than a return to pre-GDPR practices.

These firm-side adjustments provide the empirical foundation for the analysis of user privacy outcomes in Chapter 5 and economic performance and market structure in Chapter 6.

4.1 Mechanism: Risk-Based Retrenchment

This section identifies a central mechanism through which GDPR enforcement reshaped data practices in online advertising. Rather than inducing uniform reductions in data collection, the evidence shows that firms responded through what can be described as a risk-based retrenchment: selectively curtailing data practices whose expected legal and compliance risks outweighed their economic value, while preserving monetization-critical infrastructure (Miller et al., 2025; Peukert et al., 2022).

A central feature of GDPR enforcement is the expansion of joint responsibility across actors involved in data processing and sharing. Firms embedding third-party technologies or exchanging user data across organizational boundaries may be held jointly liable for downstream compliance failures, making cross-context tracking, identifier synchronization, and third-party data sharing disproportionately risky to maintain (Peukert et al., 2022; Veale & Zuiderveen Borgesius, 2022).

This mechanism operates along multiple margins. Miller et al. (2025) show that reductions in tracking technologies were concentrated among components designed to collect and share personal data across contexts, whereas monetization-critical advertising and analytics trackers were significantly less likely to be removed. Similarly, Libert, Graves, and Nielsen (2018) document that while the number of third-party domains embedded on websites changed little following GDPR enforcement, the volume and persistence of third-party cookies declined substantially.

Risk-based retrenchment was particularly pronounced in cross-firm data-sharing relationships. Urban et al. (2020) show that GDPR enforcement led firms to sever identifier-synchronization links, substantially contracting data-sharing networks without dismantling tracking infrastructure itself. Vendor choice further reflects this logic: firms simplified ad-tech stacks and shifted away

from smaller third-party vendors toward large, vertically integrated intermediaries with established compliance capacity (Johnson et al., 2023; Peukert et al., 2022).

Fixed compliance costs reinforced these patterns. Survey evidence from data-intensive startups shows that GDPR compliance required substantial organizational investment in legal expertise, documentation, and monitoring systems, disproportionately burdening smaller firms and contributing to exit from data-sharing relationships (Bessen, Impink, Reichensperger, & Seamans, 2020).

4.2 Evidence on Changes in Data Volumes, Intensity, and Flows

Consistent with risk-based retrenchment, empirical evidence documents measurable changes in data practices across multiple layers of the online advertising ecosystem. These changes affect backend data storage and processing, client-side tracking intensity, and cross-firm data sharing. Across all layers, the evidence points to selective reductions in legally risky data uses rather than a dismantling of advertising infrastructure as such.

4.2.1 Backend Data Storage and Processing

This subsection reviews evidence on how GDPR enforcement decreased firms' backend data storage and processing. Direct evidence on internal data usage is provided by Demirer, Jimenez Hernandez, Li, and Peng (2024), who study backend data storage and computation using cloud-infrastructure records. Exploiting GDPR enforcement as a regulatory shock in a difference-in-differences design, they find a 26% reduction in data storage and a 15% reduction in data processing among EU-exposed firms relative to U.S. controls. These backend reductions are directly relevant to advertising performance because modern programmatic advertising relies on massive cloud-based processing to optimize auctions and train targeting algorithms (Dubé et al., 2025). By treating data minimization as a "tax" on data intensity, the regulation forces firms to reduce the temporal depth of user profiles, constraining their ability to retain and exploit the long behavioral histories necessary for precise behavioral targeting (Demirer et al., 2024).

Complementary survey evidence from Bessen et al. (2020) shows that nearly three quarters of data-intensive startups with EU exposure deleted previously collected data and reallocated resources toward compliance-related activities. This deletion of training data directly impairs the predictive accuracy of machine learning models used for ad targeting and conversion prediction. Together, these findings indicate that GDPR constrained firms' ability to accumulate and exploit historical user data by increasing the cost of backend storage and processing, thereby weakening longitudinal

observability even when frontend tracking technologies remain in place (Bessen et al., 2020; Demirer et al., 2024).

4.2.2 Client-Side Tracking and Identifier Availability

Measurement-based studies show that GDPR enforcement reduced the density and persistence of advertising-relevant identifiers at the point of data collection. Large-scale website scans document declines in tracking libraries following enforcement, with partial retrenchment dominating wholesale exit (Degeling et al., 2019; Peukert et al., 2022). This pattern suggests adjustment at the intensive margin of tracking, reducing the volume of data collected per user, rather than the widespread abandonment of advertising-related technologies (Lefrere, Warberg, Cheyre, Marotta, & Acquisti, 2025). Evidence from European news websites shows that while third-party vendor presence remained relatively stable, the number of third-party cookies fell by approximately 22%, indicating selective reductions in identifier persistence rather than a dismantling of advertising infrastructure (Libert et al., 2018). The divergence between stable vendor presence and declining identifier counts highlights a key feature of post-GDPR adjustment: firms preserved access to advertising intermediaries while limiting the continuity and richness of user-level identification. Longitudinal data from WhoTracks.me confirm a 14.8% decline in trackers per publisher (Miller et al., 2025). Crucially, this decline was concentrated among privacy-invasive components designed for cross-context data sharing, whereas core advertising trackers remained largely resilient (Miller et al., 2025). At the extensive margin, GDPR also reduced identifier availability by inducing user opt-out. Studying an online travel intermediary, Aridor et al. (2023) document a 12.5% decline in unique cookies, reflecting users exiting client-side tracking entirely. Together, intensive-margin reductions in identifier persistence and extensive-margin opt-out jointly reduced the share of user behavior that could be reliably observed and linked over time.

4.2.3 Cross-Firm Data Sharing and Network Structure

This subsection reviews evidence on how GDPR enforcement disrupted cross-firm data-sharing networks in online advertising, leading to selective reductions in identifier synchronization and a reconfiguration toward more centralized network structures. Prior to the GDPR, programmatic advertising relied heavily on identifier synchronization to enable cross-firm data aggregation. Pre-enforcement measurements document synchronization rates of 30–40% across tracking identifiers (Acar et al., 2014; J. Wang et al., 2017), enabling widespread circulation of user-level data even when individual firms observed users only intermittently.

GDPR enforcement substantially disrupted these data-sharing networks. Urban et al. (2020) document an approximately 40% reduction in identifier-synchronization connections around enforcement, driven by firms severing data-sharing relationships rather than removing tracking technologies. While overall connectivity declined, remaining networks became more centralized, exhibiting star-like structures dominated by a small number of hubs. This reorganization implies that reductions in data sharing were unevenly distributed, disproportionately affecting peripheral intermediaries while leaving central platforms comparatively less constrained (Johnson et al., 2023; Peukert et al., 2022).

These disruptions were strongest immediately following enforcement and partially attenuated over time. Studies document modest rebounds in vendor usage and synchronization links as firms adapted compliance strategies (Johnson et al., 2023; Peukert et al., 2022; Urban et al., 2020). Importantly, this recovery reflects reconfiguration rather than a return to pre-GDPR practices: post-GDPR data-sharing architectures remain more selective and centralized, with fewer synchronization links and simplified consent chains.

Observed reductions in cross-firm data sharing should be interpreted with caution, as client-side measurements may understate total data flows if firms shift toward server-side integration or increased reliance on first-party data (Aridor et al., 2023; Urban et al., 2020). Nevertheless, the available evidence indicates a sustained reduction in observable data circulation across organizational boundaries, particularly among intermediaries lacking direct user relationships or substantial first-party data assets (Aridor et al., 2023; Johnson et al., 2023).

Taken together, the evidence shows that GDPR enforcement reshaped advertising data practices through selective retrenchment rather than uniform contraction. Firms reduced data collection, retention, and sharing where legal and compliance risks were highest, while preserving infrastructure critical to advertising monetization. Initial contractions in data flows were followed by partial adaptation and reconfiguration, resulting in a post-GDPR ecosystem characterized by lower average observability, simplified data-sharing relationships, and greater reliance on first-party data and large intermediaries.

These changes form the empirical basis for the analysis of privacy outcomes and economic performance in subsequent chapters.

5 Effects of GDPR on User Privacy

This chapter examines how the GDPR-induced changes in data practices documented in Chapter 4 translated into user privacy outcomes in practice, shifting the focus from firm-side data practices to privacy as experienced by individual users.

5.1 Privacy as Reduced Observability

This section examines how GDPR enforcement altered users' visibility within advertising data systems. Building on the concept of observability introduced in Chapter 3, privacy gains are assessed in terms of reductions in firms' ability to continuously track, link, and retain information about individual users across contexts and over time (Aridor et al., 2023; Goldberg et al., 2024).

A central implication of the changes documented in Chapter 4 is that users have become less continuously observable on average. Reductions in backend data storage and processing limit the temporal scope of user profiles (Demirer et al., 2024), while declines in persistent client-side identifiers constrain firms' ability to link behavior across visits, devices, and contexts (Aridor et al., 2023; Libert et al., 2018). Reduced identifier synchronization further weakens cross-site and cross-platform observability by limiting the construction of unified behavioral profiles spanning multiple domains (Urban et al., 2020). Together, these changes restrict both the breadth of user observation across contexts and the depth of observation over time.

On the extensive margin, increased non-consent and opt-out behavior remove a subset of users from tracking systems entirely. As discussed in Chapter 3.2, this generates a recording wedge between underlying economic activity and observed data. From a privacy perspective, the salient outcome is that some users effectively disappear from firms' data environments (Aridor et al., 2023; Goldberg et al., 2024). This form of privacy gain arises mechanically from selection rather than from changes in how observable users are tracked.

On the intensive margin, however, changes in observability are more nuanced. While the breadth of data collection declined due to reductions in third-party cookies and cross-firm identifier sharing (Libert et al., 2018; Urban et al., 2020), the persistence of observation for consenting users often increased. Aridor et al. (2023) show that although the number of observed users fell, the remaining users became trackable for longer periods of time. The mechanism is compositional. Privacy-conscious users, who previously generated noise by frequently deleting cookies or clearing identifiers, opted out of tracking altogether, leaving behind a pool of users with more stable

identifiers. As a result, firms observe fewer users overall, but those they continue to observe may be visible with greater temporal depth and predictability within specific contexts.

Consent and opt-out decisions vary systematically with privacy preferences and perceived data sensitivity, implying that some users become substantially less visible while others remain largely observable (Aridor et al., 2023; Das Chaudhury & Choe, 2023; Lin, 2022). GDPR enforcement thus reduces average observability while reshaping the composition of the observable user population. This heterogeneity is central for evaluating both the distribution of privacy gains and the informational externalities created by consent-based data regimes, as examined in Section 5.3.

5.2 Limits of Control and Consent

This section assesses whether the GDPR’s formal tools for user empowerment, such as enhanced transparency obligations, consent requirements, and data subject rights, translate into meaningful control over personal data in practice. The empirical literature suggests a systematic gap between formal choice and effective agency, driven by three frictions: information that is disclosed but not usable, consent interfaces that steer or burden rejection, and technical constraints that limit the enforceability of downstream rights in complex advertising data environments (Degeling et al., 2019; Nokhbeh Zaeem & Barber, 2020; Nouwens, Llicardi, Veale, Karger, & Kagal, 2020; Smith et al., 2024).

5.2.1 Transparency Without Comprehension

GDPR enforcement increased compliance with formal transparency requirements. Using monthly scans of the 500 most visited websites in each EU member state, Degeling et al. (2019) document a sharp rise in cookie consent banners and privacy policies around the regulation’s introduction: between January and May 2018, the share of websites displaying cookie consent notices increased by roughly 16 percentage points and the share displaying privacy policies rose by 4.9 percentage points. However, greater disclosure did not translate into meaningful comprehension. Combining automated text analysis with manual reviews of 450 privacy policies, Nokhbeh Zaeem and Barber (2020) find that many policies remain vague, incomplete, or difficult to interpret, with systematic omissions regarding data retention and sharing. Transparency therefore often increased the availability of information without improving its usability for privacy-relevant decision-making (Degeling et al., 2019; Nokhbeh Zaeem & Barber, 2020).

5.2.2 Consent Design and Behavioral Constraints

Even when consent mechanisms exist, their effectiveness depends on interface design and the costs of exercising choice. Scraping studies of consent management platforms document widespread use of choice architectures that favor acceptance, including visually salient "accept" buttons and asymmetric effort costs between acceptance and rejection (Nouwens et al., 2020; Utz, Degeling, Fahl, Schaub, & Holz, 2019). For example, Nouwens et al. (2020) show that removing a visible opt-out option from the initial interface increases acceptance rates by 22–23 percentage points.

Field evidence confirms that interaction costs are pivotal. Jha, Trevisan, Mellia, Fernandez, and Irarrazaval (2025) find that more than 60% of users reject tracking when offered a one-click "Reject All" option, whereas acceptance rates rise to approximately 90% when opting out requires additional steps. These patterns are consistent with consent fatigue: repeated requests across websites encourage users to accept tracking as a low-cost way to proceed rather than as an informed privacy choice (Dubé et al., 2025; Lefrere et al., 2025). More broadly, even users with strong stated privacy preferences frequently disclose data in response to small incentives or frictions (Athey, Catalini, & Tucker, 2017).

5.2.3 Technical Limits to Enforcing Data Rights

Beyond consent, the GDPR grants rights to access, correct, and delete personal data, but enforcing these rights depends on firms' ability to operationalize them within complex data infrastructures. Bessen et al. (2020) show that compliance required substantial organizational and technical restructuring, yet firms still face constraints when attempting to locate, trace, and delete data embedded in legacy systems, derived datasets, or machine-learning models.

These limits are particularly pronounced in online advertising, where personal data propagate across organizational boundaries. Urban et al. (2020) document that embedding a single third-party tracker can result in identifiers being synchronized with hundreds of downstream partners, complicating deletion across the supply chain. Reflecting these difficulties, Nokhbeh Zaeem and Barber (2020) note that the GDPR grants only a limited right of erasure rather than an absolute right to be forgotten. Empirical attempts to exercise these rights illustrate their limits: Subject Access Requests are often burdensome and users typically lack effective means to verify that deletion or anonymization has occurred (Acquisti et al., 2016; Urban et al., 2020).

5.2.4 Implications for User Control

Taken together, the evidence suggests that the GDPR expanded formal user rights without fully resolving the practical barriers to exercising them. Information is frequently disclosed in ways users cannot readily evaluate, consent choices are strongly shaped by interface design and interaction costs, and downstream rights such as access and erasure are difficult to enforce in multi-actor advertising data environments (Degeling et al., 2019; Nokhbeh Zaeem & Barber, 2020; Nouwens et al., 2020; Urban et al., 2020).

Moreover, even expressed choices may not reliably constrain processing in practice. Smith et al. (2024) document that 72.8% of websites using standard consent platforms continued to claim a legal basis for processing personal data even after users explicitly declined consent. This highlights a direct gap between formal consent signals and effective control over downstream data use.

5.3 Heterogeneous Privacy Effects and Externalities

The privacy effects of the GDPR are unevenly distributed across users rather than uniform. Differences in privacy preferences, resources, digital literacy, and reliance on data-driven services, shape how individuals experience GDPR-induced reductions in observability. As a result, the changes in data practices documented in Chapter 4 translate into heterogeneous privacy outcomes and externalities rather than a common privacy gain. This section examines three mechanisms behind this heterogeneity: unequal ability to exercise privacy protections, sorting induced by coarse consent regimes, and the emergence of data scarcity for already disadvantaged populations. Together, these mechanisms redistribute observability and user welfare rather than uniformly enhancing privacy.

5.3.1 Privacy Sensitivity and Unequal Protection

Although GDPR rights are formally uniform, consumers differ substantially in both their privacy preferences and their ability to act on them. A growing empirical literature suggests that privacy behaves like a luxury good: wealthier, more educated, and older consumers assign systematically higher value to protecting personal data (Dubé et al., 2025; Lin, 2022; Lin & Strulov-Shlain, 2025). Lower-income and less-educated users, by contrast, are more sensitive to defaults, framing, and interface design, making them less likely to effectively resist data collection under typical consent architectures (Lin & Strulov-Shlain, 2025).

More broadly, these patterns suggest that privacy protections under the GDPR accrue disproportionately to users with greater resources and digital literacy. When privacy requires navigating complex consent interfaces or incurring repeated interaction costs, users with higher education, technical familiarity, or stronger intrinsic privacy preferences are better positioned to effectively limit data collection. In this sense, realized privacy protection behaves like a luxury good: although legal rights are formally uniform, the ability to exercise them varies systematically across socio-economic groups (Dubé et al., 2025).

Importantly, unequal privacy protection does not imply uniformly lower data sharing among higher-status users. Instead, data sharing becomes more selective: high-income or high-engagement users are more likely to consent in high-benefit, trusted first-party contexts while avoiding low-value or opaque tracking, and technically sophisticated users are more likely to exit tracking altogether when expected benefits are low (Aridor et al., 2023; Godinho de Matos & Adjerid, 2022; Johnson et al., 2020). As a result, GDPR-induced reductions in observability are concentrated in low-benefit contexts, while observability persists where users perceive clear value.

From a welfare perspective, these patterns imply redistribution rather than uniform improvement. Privacy regulation benefits users with high intrinsic privacy costs by reducing unwanted exposure to data collection, but can lower surplus for users who value data-driven services by increasing transaction frictions and limiting the availability or discoverability of niche offerings (Campbell, Goldfarb, & Tucker, 2015). Empirical evidence supports this channel: GDPR compliance has been shown to raise search costs by increasing the time and effort required for users to locate products or information, effectively shifting welfare away from convenience and variety toward privacy protection (Zhao, Yildirim, & Chintagunta, 2023). The net effect is not a uniform privacy gain, but a reallocation of privacy and welfare across users with heterogeneous preferences and constraints.

5.3.2 Coarse Consent, Sorting, and Privacy Externalities

Heterogeneous privacy outcomes are amplified by the effective coarseness of consent choices. Experimental evidence shows substantial dispersion in willingness to share across data attributes, from negligible compensation for non-sensitive information to several dollars for sensitive data (Das Chaudhury & Choe, 2023; Lin, 2022). In practice, however, consent design and interaction costs render fine-grained preference expression rare, so that consent outcomes approximate a binary accept-or-reject decision driven by friction rather than valuation (Jha et al., 2025).

This coarseness induces sorting based on users' tolerance for interaction costs. Users unwilling to incur the cost of rejecting tracking remain broadly observable, while more privacy-sensitive users often opt out entirely. Evidence from symmetric, low-effort consent environments illustrates the strength of this effect: under Apple's App Tracking Transparency framework, consent rates fall to approximately 4% (Das Chaudhury & Choe, 2023; Dubé et al., 2025).

Sorting reshapes the composition of the observable population and generates privacy externalities. When privacy-sensitive users exit tracking systems, remaining users become more consistently identifiable over time, with repeated identifier observation increasing by about 8% (Aridor et al., 2023). Privacy gains for opt-outs therefore coexist with increased predictability for users who remain observable, implying that coarse consent regimes redistribute observability rather than uniformly reducing it.

5.3.3 Data Deserts and Exclusion Trade-Offs

While reduced observability enhances privacy for some users, it can render others digitally invisible. When personal data are scarce, fragmented, or outdated, individuals may be excluded from data-driven systems that allocate access to economic opportunities and public services.

Evidence from commercial data brokers shows substantial heterogeneity in data coverage across socio-economic groups. Low-income individuals are disproportionately likely to reside in "data deserts" characterized by missing or obsolete records across multiple data sources (Dubé et al., 2025; Neumann, Tucker, Kaplan, Mislove, & Sapiezynski, 2024; Tucker, 2023). As a result, databases used for targeting, screening, or personalization are systematically less informative for already disadvantaged populations.

Although reduced data collection lowers exposure to surveillance, it can also induce algorithmic exclusion. When firms or public institutions rely on data-driven systems to allocate credit, insurance, employment, or benefits, individuals with insufficient or unreliable data are often excluded because their characteristics cannot be confidently inferred (Dubé et al., 2025). GDPR-induced reductions in data collection therefore enhance privacy for some users while increasing exclusion risks for others, reinforcing the uneven distribution of observability and welfare rather than uniformly improving outcomes (Dubé et al., 2025; Tucker, 2023).

6 Effects of GDPR on Efficiency and Firm Performance

This chapter reviews empirical evidence on how GDPR-induced limits on personal data collection affect advertising effectiveness, firm performance, and market structure in online advertising markets. The literature identifies effects along three dimensions of economic efficiency: allocative efficiency, through constraints on targeting and measurement in advertising auctions; productive efficiency, via compliance costs and changes in firms' access to data; and dynamic efficiency, by reshaping incentives for entry, innovation and scale in data-intensive markets (as discussed in Chapter 3.4).

Across these dimensions, GDPR operates through two channels: reduced availability of user-level data and changes in the scale and composition of observable users due to consent and opt-out behavior. The following sections examine how these channels propagate from advertising performance to firm outcomes and market structure.

6.1 Direct Effects on Advertising Effectiveness

This section reviews empirical evidence on how GDPR enforcement affected digital advertising effectiveness, focusing on allocative outcomes such as click-through rates, conversions, bids, and advertising revenue. By restricting access to persistent identifiers and cross-context behavioral data, GDPR reduced advertisers' ability to target and measure users, leading to economically meaningful declines in effectiveness.

In response, advertisers may substitute toward contextual targeting, which conditions ad placement on page content rather than user histories (Zhang & Katona, 2012) and is therefore less affected by consent constraints, albeit less precise than behavioral targeting (P. Wang et al., 2024).

6.1.1 Baseline Value of Data-Driven Targeting

To gauge the magnitude of post-GDPR performance changes, it is useful to establish the baseline value of behavioral targeting. A large pre-regulation literature shows that access to user-level behavioral data substantially improves advertising effectiveness. Early evidence indicates that behavioral targeting increases click-through rates by several hundred percent relative to non-targeted ads (Yan et al., 2009), while more recent work shows that machine-learning-based targeting raises average click-through rates by about 67% compared to traditional segmentation (Rafieian & Yoganarasimhan, 2021). At large platforms, eliminating access to personal data reduces click-through rates by up to 75% and purchases by more than 80% (Sun, Yuan, Li, Zhang,

& Xu, 2024). It is important to note that estimates of the value of behavioral targeting may overstate its true causal impact. Farahat and Bailey (2012) show that a substantial share of the measured effectiveness of targeted advertising reflects selection effects: users who are targeted are already more likely to convert even without seeing an ad. As a result, the loss of targeting data following GDPR enforcement removes some measured effectiveness that reflects selection rather than eliminating all underlying economic value. This helps explain why post-GDPR revenue declines are economically meaningful but generally more modest than pre-regulation estimates of targeting gains would suggest (P. Wang et al., 2024).

Against this backdrop, even modest post-GDPR declines represent economically meaningful losses in allocative efficiency. Because advertising markets are highly optimized, marginal reductions in targeting accuracy and attribution can translate into significant revenue effects (Goldberg et al., 2024; P. Wang et al., 2024). GDPR therefore constitutes a negative shock to an otherwise technologically advanced advertising environment rather than to an inefficient baseline (Rafieian & Yoganarasimhan, 2021).

6.1.2 Evidence from Publisher-Side Data

Using proprietary data from a large U.S.-based publisher operating under a pay-per-click model, P. Wang et al. (2024) provide causal evidence on the effects of GDPR compliance in a setting that avoids data-collection bias from intermediary non-consent. Because performance, bidding, and revenue outcomes are recorded directly by the publisher, advertising effects are observed even for users who do not consent to tracking (P. Wang et al., 2024).

Exploiting a difference-in-differences design comparing EU and non-EU IP addresses, the authors isolate the loss of personalization while holding ad exposure constant. Unlike platforms that applied GDPR globally, this publisher requested consent only from EU users, yielding a credible non-EU control group. GDPR compliance led to statistically significant declines in advertising performance: click-through rates fell by 2.1%, conversion rates by 5.4%, and revenue per click by 5.7%, driven by lower bid prices and reduced advertiser participation (P. Wang et al., 2024).

These effects are heterogeneous across advertisers. Performance declines were largest in data-intensive sectors such as travel and financial services and substantially smaller in retail and consumer packaged goods, reflecting differential reliance on precise targeting. As a result, GDPR disproportionately reduced demand from high-value advertisers, amplifying publisher revenue losses (P. Wang et al., 2024).

These estimates likely represent a lower bound. The publisher possesses strong first-party assets and substantial capabilities for contextual targeting. With about 47% of impressions appearing on thematically specific pages and relatively high opt-in rates under a uniform consent strategy, these findings imply larger negative effects for smaller or more intermediary-dependent publishers (P. Wang et al., 2024).

6.1.3 Advertising as a Traffic-Generation Mechanism

Complementary evidence indicates that GDPR also weakened advertising as a mechanism for directing users to publisher websites. Congiu et al. (2022) analyze traffic and engagement data for approximately 5,000 web domains, decomposing visits by acquisition channel. They find that paid traffic channels experienced the largest and most persistent declines following GDPR enforcement. In the long run, traffic from display advertising fell by approximately 38% and was accompanied by declines in user engagement, including shorter visit durations and fewer pages viewed per session (Congiu et al., 2022).

This finding is corroborated by Goldberg et al. (2024) using proprietary website analytics from over 1,000 firms. They document that following GDPR enforcement, total recorded pageviews fell by 11.7% among EU users. Crucially, this decline was not uniform: traffic originating specifically from display advertising experienced a significantly larger drop of approximately 26.5% compared to direct traffic or search. These patterns across both studies are consistent with reduced match quality between ads and users in the absence of detailed behavioral data. Although neither study observes auction outcomes directly, their results complement auction-level evidence by showing that GDPR weakened advertising's effectiveness as a matching and discovery technology for the broader web ecosystem.

6.1.4 Evidence from Intermediary Data

Evidence from advertising intermediaries highlights how GDPR affects advertising effectiveness in environments where non-consenting users are excluded entirely from ad exposure. Aridor et al. (2023) exploit GDPR as a natural experiment using data from a travel-sector online advertising intermediary operating under RTB. Following GDPR enforcement, advertising scale contracted sharply: total cookies fell by 12.5% and ad clicks by 13.5%, mechanically reducing revenue due to diminished observability (Aridor et al., 2023).

Over time, however, prices adjusted. Although average bids were initially unchanged, they increased by roughly 12% for remaining users. The authors attribute this to "privacy means

substitution”: the exit of users generating measurement noise left a smaller but more consistently identifiable user pool. Advertisers valued this improved attribution, reflecting an improvement in observability on the intensive margin, as discussed in Section 5.1. Higher bids partially offset the contraction in advertising scale, so that despite a permanent loss of over 12% of users, the post-GDPR revenue decline was ultimately not statistically significant (Aridor et al., 2023).

6.1.5 Reconciling Evidence Across Institutional Settings

Taken together, the evidence highlights two complementary channels through which GDPR affects advertising effectiveness. In settings where non-consenting users continue to receive ads, GDPR weakens personalization and targeting. P. Wang et al. (2024) show that reduced match quality lowers click-through rates and conversions, leading advertisers to cut bids and participation.

By contrast, in environments where opt-out users are excluded entirely from ad exposure, GDPR primarily reduces advertising scale but can improve attribution among remaining users. Aridor et al. (2023) document a contraction in advertising volume alongside an approximately 12% increase in average bids, reflecting a composition effect: privacy-sensitive users, who previously generated measurement noise by deleting cookies, exit the market, leaving a smaller but more consistently trackable user pool.

These patterns reflect differences in institutional design and data architecture rather than conflicting findings. Advertiser bids depend on both expected user value (targeting) and confidence in measurement (attribution). Restrictions on persistent identifiers reduce expected value by limiting behavioral targeting (Goldfarb & Tucker, 2011; P. Wang et al., 2024), but can increase confidence by eliminating noisy or fragmented identifiers (Aridor et al., 2023). Publisher revenue ultimately depends on whether attribution-driven price premiums offset volume losses from user opt-outs.

6.1.6 Contextual Targeting and Selection Effects

While advertisers substitute toward contextual targeting, this mechanism offers only partial compensation for the loss of behavioral data. Using publisher-side data, P. Wang et al. (2024) show that relevant page context recovers about 44% of the conversion-rate decline caused by reduced access to personal data. This indicates that while contextual signals are valuable, they are imperfect substitutes for individual-level tracking.

This performance gap aligns with prior literature showing that while contextual and behavioral information jointly maximize performance (Lu, Zhao, & Xue, 2016), behavioral data remain

substantially more predictive of clicks and conversions than context alone (Rafieian & Yoganarasimhan, 2021). Consequently, substitution toward contextual targeting mitigates but does not eliminate the allocative efficiency losses arising from restrictions on persistent identifiers (Johnson et al., 2020).

6.2 Effects on Firm Performance and Market Structure

The changes in advertising performance documented in Section 6.1 have implications that extend beyond individual auctions and campaigns. This section shifts the focus from advertising effectiveness to firm-level outcomes, examining how GDPR-induced changes in data availability, compliance costs, and advertiser participation translate into differences in revenue, profitability, and competitive dynamics across firms.

6.2.1 From Advertising Performance to Firm-Level Revenue and Profitability

GDPR-induced changes in advertising performance translate into economically meaningful revenue effects. In the study introduced in Section 6.1, P. Wang et al. (2024) show that GDPR compliance reduced revenue per click by 5.7% for the studied publisher, reflecting a true economic loss driven by lower bids and reduced advertiser participation rather than reduced observability.

In their study on website analytics, introduced earlier (see Section 6.1: Advertising as a Traffic-Generation Mechanism), Goldberg et al. (2024) document a 13.3% decline in recorded e-commerce revenue following GDPR enforcement, with effects nearly twice as large for smaller, advertising-dependent firms. The discrepancy highlights the recording wedge, introduced in Section 3.2. Outcomes in Goldberg et al. (2024) mechanically decline when non-consenting users become unobserved, whereas P. Wang et al. (2024) observe publisher-side revenue directly and continue to capture untargeted monetization. The gap between the roughly 13% decline in recorded revenue and the 5.7% decline in revenue per click is therefore consistent with reduced observability accounting for a substantial share of the apparent revenue loss.

Consistent with these revenue effects, Frey and Presidente (2024) estimate that GDPR reduced the profits of technology firms serving EU markets by about 2.1%, with disproportionately larger losses among smaller and more data-intensive firms, reflecting both higher compliance costs and diminished returns to data-intensive inputs such as targeted advertising (Chen, 2022; Frey & Presidente, 2024).

GDPR effects are heterogeneous across firms. While advertising-dependent intermediaries experience revenue declines, firms with strong first-party relationships may benefit. Godinho de Matos and Adjerid (2022) show that GDPR-style re-elicitation increased consent among high-spending, long-tenure customers for a large telecommunications firm, raising targeted marketing effectiveness by 16% and sales by 3.7%. These gains reflect improved trust and data quality in first-party settings, indicating that GDPR reallocates advertising effectiveness toward firms with direct consumer relationships. While much empirical evidence points to efficiency losses, recent theory shows that opt-in regimes can increase welfare when firm revenues are primarily usage-based rather than data-based (Choe et al., 2025).

6.2.2 Differential Impacts and Scale Effects

GDPR-induced shocks are uneven across firms. Reductions in data availability, higher effective costs of data use, and largely fixed compliance investments disproportionately affect advertising-dependent firms, while larger incumbents are better positioned to absorb these constraints due to scale, diversification, and access to first-party data.

As shown in the previous subsection, Goldberg et al. (2024) find that revenue losses among small e-commerce firms are nearly twice as large as among larger firms following GDPR enforcement. On the advertiser side, P. Wang et al. (2024) document disproportionately large liquidity reductions in data-intensive niche sectors such as travel, where precise targeting is critical. Advertiser exit in these sectors generates a negative demand shock that further disadvantages smaller publishers lacking the diversified advertiser base of large platforms.

A key mechanism is an increase in the effective cost of using personal data. Demirer et al. (2024) estimate that GDPR raises the marginal cost of personal data usage by roughly 20%, acting as a "data tax" that reduces data storage and processing intensity. Because compliance costs are largely fixed, larger firms can amortize them more easily and rely on stronger first-party data assets, consistent with predictions that regulation favors generalist firms over smaller specialists (Bessen et al., 2020; Campbell et al., 2015).

Evidence on size effects, however, is not uniform. While Goldberg et al. (2024) and Johnson et al. (2023) find a monotonic protective effect of size, Congiu et al. (2022) document an inverted U-shaped relationship between website size and post-GDPR traffic losses. In their study of 5,000 domains, both small and very large sites experience the largest declines, while medium-sized firms are relatively insulated. This pattern suggests that scale confers compliance advantages (Campbell

et al., 2015) but also attracts greater regulatory and consumer scrutiny, implying that GDPR's net effect depends on whether efficiency gains from scale outweigh the costs of being a primary target for privacy-conscious users (Congiu et al., 2022).

6.2.3 Content Supply and Output Responses

This subsection examines whether GDPR-induced revenue losses translated into real-side adjustments in content supply and output among advertising-funded publishers. Despite documented declines in advertising revenue following GDPR enforcement across e-commerce and ad-supported digital markets (Goldberg et al., 2024; P. Wang et al., 2024), available evidence suggests limited adjustment on the content supply margin in the short run. Studying publisher behavior after GDPR enforcement, Lefrere et al. (2025) find no significant reductions in content volume or observable quality proxies among news and media websites. Rather than contracting output, publishers appear to have absorbed revenue shocks through changes in monetization technology, inventory composition, and the addition of compliance-related fixed costs.

This pattern is consistent with adjustment along margins other than content provision. Digital content production is characterized by high fixed costs and low marginal costs, making short-run reductions in output unattractive even when revenues decline. Moreover, advertising revenue often finances content through cross-subsidization across users and formats, allowing publishers to smooth revenue shocks by rationalizing their tracking stacks and adopting consent frameworks rather than immediately reducing content supply (Lefrere et al., 2025). As a result, GDPR-induced declines in advertising effectiveness primarily affected monetization and intermediation rather than the real-side production of digital content in the short run (Lefrere et al., 2025).

However, the ability to sustain output was uneven across firms. Larger and medium-sized publishers with diversified revenue streams, established brands, or strong first-party relationships were better positioned to absorb compliance costs and revenue losses. Goldberg et al. (2024) find that revenue declines were nearly twice as large for small websites than for larger ones, indicating that aggregate content stability may mask compositional change. Similarly, Congiu et al. (2022) document an inverted U-shaped relationship between website size and post-GDPR traffic outcomes: small and very large sites experienced traffic losses, while medium-sized, trusted publishers were relatively insulated. These patterns suggest that content resilience among incumbents may coexist with tighter constraints for smaller, advertising-dependent firms.

It is also important to recognize the limits of available measures. The finding of content resilience in Lefrere et al. (2025) relies on automated natural language processing techniques to construct observable quality proxies, such as article length, sentence complexity, and standardized readability scores. While these measures capture salient structural features of published content, they are imperfect proxies for users' experienced quality or the editorial depth of reporting. Stability in these metrics does not rule out more subtle adjustments, such as reductions in investigative investment, topic diversity, or editorial risk-taking, which are margins that automated tools cannot readily detect.

Evidence from the mobile app ecosystem, another advertising-dependent sector, underscores the relevance of such dynamic margins. Following GDPR enforcement, app markets experienced a sharp contraction in entry, effectively generating a "lost generation" of new products. Using comprehensive Google Play Store data, Janßen et al. (2022) show that new app entry declined by 47.2% after GDPR enactment, with effects concentrated among small and advertising-reliant developers. Although these findings do not directly measure web content supply, they suggest that privacy-induced monetization constraints may disproportionately affect experimentation and innovation even when observable output among incumbents remains stable.

6.2.4 Market Structure and Concentration

Beyond firm-level heterogeneity, GDPR reshaped market structure by raising the cost of operating complex, multi-vendor advertising stacks. Tighter compliance requirements and limits on cross-firm data sharing encouraged publishers to simplify ad-tech arrangements and reallocate activity toward large, vertically integrated platforms. Firm-level evidence shows substantial contraction and consolidation. Johnson et al. (2023) find that websites serving EU users reduced third-party advertising and analytics vendors by about 15%, disproportionately affecting smaller specialists. This contraction is driven by GDPR's joint liability provisions, which expose publishers to partners' compliance failures and incentivize vendor pruning in favor of established incumbents (Peukert et al., 2022). Using data from over 110,000 websites, Peukert et al. (2022) show that while overall third-party usage declined, market shares became more concentrated, with Google gaining share in advertising and analytics despite market contraction.

Network-level evidence indicates reorganization rather than uniform centralization. Urban et al. (2020) document a roughly 40% decline in identifier synchronization after GDPR, while remaining data-sharing networks retained a star-like topology dominated by a few hubs. Although Google maintained extensive reach, its betweenness centrality fell by 36%, consistent with a shift toward

more closed first-party ecosystems. By contrast, Oracle's betweenness centrality rose by 71%, reflecting backend reconfiguration as networks fragmented. Consistent with this pattern, Peukert et al. (2022) show that despite Google's category-level gains, the number of websites serving Google requests fell by about 2.4%, indicating absolute scale losses even for dominant firms.

Evidence also points to partial but asymmetric recovery. Johnson et al. (2023) document a rebound in third-party vendor usage, and Urban et al. (2020) find modest recovery in identifier synchronization. Re-entry, however, was concentrated among large, vertically integrated platforms, while smaller and specialized vendors were less likely to return, implying consolidation rather than restoration of pre-GDPR fragmentation.

Finally, GDPR spurred rapid growth in privacy compliance intermediaries. Johnson et al. (2023) show that "Privacy Compliance", a category dominated by CMPs, was the only vendor category to expand immediately after enforcement, growing by over 20% in the short run and more than 100% by the end of 2018. Peukert et al. (2022) similarly identify Consent Management Platforms among the largest market-share gainers, indicating that CMPs emerged as a gatekeeping layer that raises fixed participation costs and reinforces concentration.

6.2.5 Entry, Innovation, and Dynamic Efficiency

This subsection examines how GDPR-induced data constraints affected entry, innovation, and dynamic efficiency in data-intensive markets.

Evidence from venture capital markets suggests that GDPR is associated with reduced early-stage entry. Using state-level panel data, Jia et al. (2021) estimate that EU technology ventures experienced an average decline of \$3.38 million in weekly venture capital raised per state relative to comparable U.S. regions following enforcement. This decline is concentrated among young ventures aged zero to three years, while more mature firms are largely unaffected, indicating reduced early-stage experimentation.

Stronger dynamic effects appear in app markets, where monetization is highly advertising-dependent. As discussed before, Janßen et al. (2022) show that GDPR enforcement led to the exit of over one-third of available apps and reduced new entry by 47.2%. They further estimate a long-run consumer surplus reduction of about 32%, driven primarily by reduced variety and innovation rather than lower usage of surviving apps.

These patterns are consistent with theoretical mechanisms emphasizing the role of data access in innovation. Hagiu and Wright (2023) distinguish between data stocks, as in accumulated historical data, and data flows, describing the continuous generation of new data. Incumbent firms possess substantial data stocks that allow them to maintain product quality and algorithmic performance even when new data collection is constrained. By contrast, entrants depend critically on data flows to train models, experiment, and refine products. By restricting these flows through consent requirements and limits on data collection, privacy regulation disproportionately raises entry costs for new firms while insulating incumbents.

Complementary evidence supports this mechanism. Janßen et al. (2022) show that apps exiting after GDPR exhibited fewer feature updates even prior to enforcement, suggesting limited capacity to absorb additional data constraints. Survey evidence similarly indicates that compliance costs led startups to abandon products and reallocate resources away from innovation (Bessen et al., 2020). This divergence helps reconcile the stability of incumbent content supply documented by Lefrere et al. (2025) with the collapse in new entry observed in venture capital and app markets. While established firms were able to adapt monetization and data practices to preserve output, restrictions on data flows curtailed experimentation and entry, generating substantial losses in dynamic efficiency.

6.3 Actor-Specific Adjustments and Implications

The firm-level and market-structure effects documented in Section 6.2 propagated unevenly through the digital advertising ecosystem, prompting distinct strategic responses by publishers, advertisers, and intermediaries. Rather than uniformly scaling back activity, actors adapted by reconfiguring compliance, reallocating risk, and concentrating operations in environments where data use remained legally and economically viable. These adjustments reflect differences in exposure to observability losses, fixed compliance costs, and reliance on cross-firm data sharing.

6.3.1 Publishers: Compliance Layering and Risk Reallocation

Publishers faced the immediate challenge of legitimizing advertising-based monetization under heightened legal exposure. Their primary response was to layer compliance onto existing monetization systems rather than dismantling them. A central component of this strategy was the widespread adoption of consent management platforms (CMPs), which standardized consent interfaces and mediated access to user data for downstream vendors (Hils et al., 2020). CMPs

allowed publishers to preserve access to advertising demand while shifting legal responsibility and operational complexity to a dedicated compliance layer.

Publishers also segmented compliance geographically, applying consent prompts and tracking restrictions primarily to EU users while maintaining less constrained practices for non-EU traffic (P. Wang et al., 2024). In parallel, they pruned legally risky third-party vendors and consolidated around fewer, larger partners perceived as safer under joint-liability provisions (Johnson et al., 2023; Libert et al., 2018; Peukert et al., 2022; Urban et al., 2020). This consolidation reduced the breadth of data sharing while preserving core monetization channels.

These adjustments allowed publishers to preserve advertising inventory and content supply while reallocating regulatory risk within their ad-tech stacks. Adaptation capacity, however, was uneven. Firms with trusted first-party relationships were better able to sustain addressability through higher consent rates (Godinho de Matos & Adjerid, 2022), whereas smaller publishers faced tighter constraints, consistent with fixed compliance costs disproportionately burdening low-scale operations (Campbell et al., 2015). As a result, the GDPR amplified pre-existing scale advantages in publisher monetization.

6.3.2 Advertisers: Reallocation Toward Reliable Attribution

Advertisers adapted to the GDPR primarily by reallocating budgets toward environments where targeting and attribution remained reliable. As persistent identifiers became less available on the open web, advertisers reduced participation in low-information contexts and concentrated spending on channels, platforms, and user segments with higher measurement confidence. This adjustment reflects a shift along the intensive margin of advertising performance, favoring precision and measurability over broad reach (Aridor et al., 2023; Godinho de Matos & Adjerid, 2022; Johnson et al., 2023; P. Wang et al., 2024).

This reallocation favored large platforms with strong first-party data and integrated consent frameworks, whose logged-in environments preserved addressability that was lost elsewhere (Johnson et al., 2023; Peukert et al., 2022). Advertisers also substituted partially toward contextual targeting, which mitigated some losses from restricted behavioral data but remained an imperfect substitute for individual-level tracking (Rafieian & Yoganarasimhan, 2021; P. Wang et al., 2024). Contextual approaches reduced reliance on personal data but offered weaker learning and attribution signals.

Overall, advertiser responses reflect a shift away from scale-driven strategies toward environments offering higher attribution reliability, involving persistent trade-offs between audience reach, targeting precision, and measurement accuracy.

6.3.3 Intermediaries: Consolidation and Gatekeeping

Advertising-technology intermediaries were among the most heavily affected actors following GDPR enforcement, given their reliance on large-scale, cross-firm data sharing. By imposing strict consent and joint-liability requirements across the data supply chain, the regulation destabilized business models reliant on large-scale, cross-firm data sharing.

A central friction arose from the incompatibility between GDPR consent requirements and real-time bidding. Publishers cannot obtain specific, informed consent for the hundreds of downstream demand-side platforms that may participate in an auction, creating legal uncertainty (Veale & Zuiderveen Borgesius, 2022). In response, publishers pruned intermediary stacks, disproportionately removing smaller and specialized SSPs and DSPs perceived as risky (Johnson et al., 2023). This adjustment reduced intermediary participation at the extensive margin, compressing the number of active market participants.

At the data aggregation layer, reduced identifier synchronization weakened open data-sharing networks on which independent DMPs and data brokers relied (Urban et al., 2020). Rather than collapsing uniformly, aggregation reconfigured toward more closed, backend arrangements dominated by a small number of large intermediaries. These actors could internalize data flows and compliance within vertically integrated structures.

These dynamics reinforced the position of vertically integrated platforms that combine publishing, intermediation, and first-party data within unified consent frameworks (Johnson et al., 2023; Peukert et al., 2022). CMPs emerged as a complementary gatekeeping layer, mediating consent signals between users and the broader advertising ecosystem and establishing themselves as essential compliance infrastructure (Hils et al., 2020; Johnson et al., 2023). Gatekeeping thus shifted from open-market coordination toward control over compliant access to user data.

Overall, GDPR reshaped the advertising ecosystem not by eliminating data-driven advertising, but by compressing intermediation, reallocating compliance risk, and concentrating activity in environments capable of sustaining lawful data use at scale. The resulting market structure reflects lower average observability, higher fixed costs of participation, and increased concentration among actors with first-party data and compliance capacity.

7 Privacy and Performance under the GDPR: A Synthesis

The preceding chapters analyzed how the GDPR reshaped data practices, privacy outcomes, and economic performance in online advertising markets. Taken together, the evidence supports a coherent conclusion: the GDPR increased average privacy primarily by reducing user observability through compliance-driven retrenchment and user exit, but did so at the cost of allocative and dynamic efficiency, while redistributing data access, surplus, and market power across users and firms.

This chapter synthesizes these findings to answer three questions. First, did the GDPR achieve its stated privacy objectives when evaluated against its own normative benchmarks? Second, how should the documented economic effects be assessed relative to an appropriate welfare perspective? Third, how were privacy gains and efficiency costs distributed across users, firms, and segments of the digital advertising ecosystem?

7.1 Intended vs. Realized Privacy: Data Minimization without Effective User Control

A central objective of the GDPR is not merely to reduce the use of personal data, but to do so through user empowerment: limiting unnecessary collection while enabling individuals to exercise meaningful control over how their data are processed (GDPR Recital 7; Congiu et al., 2022; Godinho de Matos & Adjerid, 2022; Nokhbeh Zaeem & Barber, 2020). Evaluating whether this objective is achieved therefore requires distinguishing between privacy as an outcome, as in reduced data collection and observability (Johnson et al., 2023; Peukert et al., 2022), and privacy as a mechanism, as in active, informed, and granular user control (Aridor et al., 2023; Nouwens et al., 2020).

Evaluated as a data-minimization policy, the GDPR can be considered broadly successful. Across a wide range of settings, enforcement is associated with sustained reductions in data use and circulation, consistent with an increase in the effective cost of treating personal data as an input into advertising and analytics processes (Demirer et al., 2024; Libert et al., 2018; Miller et al., 2025; Urban et al., 2020). From this perspective, the regulation largely achieves its objective of reducing aggregate observability, even though the specific channels differ across firms and market segments (Aridor et al., 2023; Johnson et al., 2023; Miller et al., 2025; Peukert et al., 2022).

From the perspective of user empowerment, however, the GDPR performs substantially worse. Rather than functioning as a mechanism for active choice, consent largely operates as a legal

boundary condition shaping firm compliance behavior (Hils et al., 2020; Veale & Zuiderveen Borgesius, 2022). Consent interfaces frequently rely on asymmetric choice architectures that steer users toward acceptance without enabling informed or deliberative decision-making (Jha et al., 2025; Nouwens et al., 2020; Utz et al., 2019), and formal refusals are often weakly binding, as firms may continue processing personal data under alternative legal bases such as legitimate interest (Smith et al., 2024). Beyond consent, GDPR access and control rights—including rights to access, erasure, and objection—prove costly to exercise and difficult to enforce in practice, constraining their effectiveness as tools for active privacy management (Bessen et al., 2020; Nokhbeh Zaeem & Barber, 2020; Urban et al., 2020). Taken together, these frictions imply that the GDPR expands formal legal protections without generating broad-based user agency.

Observed privacy gains therefore arise primarily through channels other than active user control. Firm-side compliance-driven retrenchment and user exit from tracking systems account for the dominant share of reduced observability, while fine-grained, user-initiated data management plays a secondary and often symbolic role (Aridor et al., 2023; Demirer et al., 2024; Goldberg et al., 2024). In equilibrium, the GDPR functions less as a tool of user agency and more as a compliance framework in which choices are frequently constrained by interface design or overridden by alternative legal bases.

These outcomes are highly heterogeneous across users. Individuals who exit tracking environments or interact with privacy-protective defaults experience substantial gains, while users who remain observable often become more predictable (Aridor et al., 2023). Average privacy improvements therefore mask redistribution and externalities rather than uniform increases in individual agency. Emerging evidence further suggests that these dynamics may exacerbate inequality, favoring users with greater digital literacy or resources while deepening exclusion for populations already marginalized within digital markets (Dubé et al., 2025; Lin, 2022).

Importantly, these patterns should not be interpreted as inherent limits of consent-based regulation. When consent is implemented transparently and without manipulative defaults, it can support higher participation and increase the economic value of consenting users (Godinho de Matos & Adjerid, 2022). The gap between the GDPR's stated objectives and its realized effects thus reflects equilibrium implementation choices and strategic design responses rather than a fundamental incompatibility between data minimization and meaningful user control.

Finally, privacy gains under the GDPR extend beyond tracking and targeting. Independent of user choice, the regulation strengthens incentives for data security through breach notification

requirements, liability provisions, and enforcement risk, leading firms to increase investments in safeguarding stored data (Ke & Sudhir, 2023). These defensive privacy gains operate orthogonally to user empowerment and imply stronger protection against unauthorized access even where commercial data processing persists.

Taken together, the evidence supports a clear ranking of privacy channels under the GDPR. Firm-side retrenchment, user exit, and enhanced data security account for the bulk of observed privacy gains, while active, fine-grained user control plays a limited role. The regulation reshapes who is observed and how data are governed, but largely falls short of establishing users as active managers of their own digital footprints.

7.2 Efficiency Costs of Privacy Regulation

Privacy regulation alters the production environment of the digital economy by raising the cost of using personal data as an input. Because data are central to reducing search frictions, improving matching, and enabling measurement in advertising markets, these constraints generate efficiency costs distinct from the privacy gains discussed in Section 7.1 and the distributional consequences analyzed in Section 7.3. This section synthesizes the evidence on these costs across three margins: allocative efficiency in matching and measurement, productive efficiency in data use and compliance, and dynamic efficiency in entry, experimentation, and innovation.

Given the highly optimized nature of digital advertising markets, restrictions on data access introduced by privacy regulation generate allocative efficiency costs by degrading the precision of matching and measurement (Rafieian & Yoganarasimhan, 2021; P. Wang et al., 2024). At the same time, these costs are bounded by the extent to which pre-regulation targeting reflected selection rather than causal persuasion, implying that post-regulation performance declines represent losses in incremental allocative efficiency rather than wholesale reversals of pre-GDPR gains (Farahat & Bailey, 2012; Goldberg et al., 2024).

Allocative efficiency losses arise because reduced data availability weakens targeting and attribution, leading to systematically less efficient matches between advertisers and users even after substitution toward privacy-preserving alternatives (Goldfarb & Tucker, 2011; Johnson et al., 2020; Rafieian & Yoganarasimhan, 2021). These losses take the form of deadweight inefficiency rather than mere redistribution (P. Wang et al., 2024). Their manifestation depends on institutional design: in some settings, inefficiency appears as weaker personalization among served users, while in others it takes the form of reduced advertising scale with partial offsetting gains from improved

measurement among remaining users (Aridor et al., 2023; P. Wang et al., 2024). In both cases, allocative efficiency remains lower relative to a higher-information benchmark.

Productive efficiency is affected through increased compliance, consent-management, and governance frictions. Privacy regulation raises both fixed and marginal costs by requiring firms to devote real resources to legal risk mitigation, auditing, and data governance (Bessen et al., 2020). In addition, consent and recording requirements introduce transaction and observability frictions that prevent some data from being collected, transmitted, or used even when underlying economic activity continues (Goldberg et al., 2024). These losses reflect productive inefficiency insofar as effective input use is reduced by governance and measurement constraints, effectively a tax on data processing, rather than by deliberate privacy choices (Demirer et al., 2024).

The most consequential efficiency costs arise on dynamic margins. By restricting data flows, privacy regulation disproportionately affects entry, experimentation, and innovation (Janßen et al., 2022; Jia et al., 2021). Incumbent firms rely primarily on accumulated data stocks to sustain existing products, whereas entrants depend on continuous data flows to learn and iterate (Hagiu & Wright, 2023). Raising the effective cost of data collection therefore increases barriers to experimentation and disproportionately constrains new firms (Bessen et al., 2020; Demirer et al., 2024). Innovations that would otherwise occur may be delayed or foregone entirely, generating dynamic efficiency losses that are not immediately visible in short-run output measures but compound over time through reduced variety and slower technological progress (Jones & Tonetti, 2020).

Privacy regulation also reshapes firm boundaries in ways that have efficiency implications. By increasing transaction costs associated with consent, compliance, and joint liability, regulation encourages firms to internalize activities previously coordinated through specialized third-party providers (Peukert et al., 2022; Urban et al., 2020). Larger, multi-product firms are better able to amortize these costs, while smaller or more specialized firms face higher per-unit burdens (Campbell et al., 2015; Demirer et al., 2024). The resulting shift toward more integrated arrangements can reduce regulatory exposure but may also displace efficient specialization and modular production, lowering productive efficiency at the system level (Johnson et al., 2023; Urban et al., 2020). These effects reflect efficiency mechanisms rather than distributional outcomes, which are examined in Section 7.3.

Taken together, the evidence supports a clear ranking of efficiency costs under privacy regulation. Allocative inefficiencies in matching and measurement emerge quickly and affect incumbent firms

in the short run (Goldberg et al., 2024; P. Wang et al., 2024). Productive inefficiencies arise through compliance and governance frictions that absorb real resources (Bessen et al., 2020; Demirer et al., 2024). The largest potential welfare costs stem from dynamic inefficiency: reduced entry, experimentation, and innovation that constrain future output and variety (Hagiu & Wright, 2023; Janßen et al., 2022; Jia et al., 2021; Jones & Tonetti, 2020). These efficiency losses, however, do not mechanically translate into lower total welfare, as opt-in privacy regimes can increase welfare when firms rely primarily on usage-based rather than data-based revenues (Choe et al., 2025). This distinction helps reconcile the relative stability of incumbent performance in Chapter 6 with concerns about long-run welfare losses that are not immediately visible in aggregate output measures (Johnson et al., 2023; Peukert et al., 2022).

In sum, privacy regulation imposes real efficiency costs by weakening matching, raising production costs, and constraining innovation. While short-run effects may appear modest in static metrics, the long-run welfare implications hinge primarily on dynamic margins, where efficiency losses accumulate over time.

7.3 Redistribution across the Advertising Value Chain

The privacy gains and efficiency costs associated with the GDPR are unevenly distributed across the digital advertising ecosystem. Rather than uniformly reducing data use or economic surplus, the regulation redistributes observability, market power, and participation across users, firms, and layers of intermediation. These distributional effects arise from asymmetric exposure to compliance costs, differential access to first-party data, and variation in reliance on cross-firm data flows (Campbell et al., 2015; Gal & Aviv, 2020).

A central distributional margin concerns scale and compliance capacity. Actors able to absorb fixed compliance costs and manage legal risk internally are systematically advantaged, while smaller and more specialized participants are disproportionately constrained. Among intermediaries, GDPR enforcement induces contraction and consolidation: smaller tracking, analytics, and data vendors exit or lose scale, while large, compliance-capable platforms expand their relative position (Johnson et al., 2023; Peukert et al., 2022). At the same time, compliance-oriented intermediaries, most notably consent management platforms, benefit directly from the expansion of regulatory overhead, and emerge as essential gatekeepers to lawful monetization of EU traffic (Hils et al., 2020; Johnson et al., 2023). Data aggregation does not disappear but reconfigures, shifting away from open, cross-firm sharing toward closed first-party ecosystems and centralized backend brokers (Peukert et al.,

2022; Urban et al., 2020). The net effect is a redistribution of surplus toward actors that control compliant access to data rather than those specializing in its exchange or analysis.

A second distributional fault line reflects heterogeneity in publisher business models and data dependence. GDPR does not uniformly disadvantage publishers; instead, it reallocates advertising viability. Small, advertising-dependent publishers experience disproportionate revenue losses, while medium-sized, trusted publishers are relatively insulated, consistent with an inverted U-shaped relationship between scale and post-GDPR outcomes (Congiu et al., 2022). Large aggregators and generic platforms also face constraints linked to their reliance on user-level data. These patterns are especially pronounced in mobile app markets, where enforcement accelerates exit among small developers and compresses the long tail of available products (Janßen et al., 2022; Miller et al., 2025). Aggregate content stability therefore masks substantial redistribution within the publishing sector (Lefrere et al., 2025).

These publisher-level effects are reinforced by advertiser reallocation behavior. Advertisers respond to degraded targeting and measurement by shifting spending toward environments with stronger consent coverage, richer contextual signals, or robust first-party data (Goldberg et al., 2024; P. Wang et al., 2024). This reallocation favors large platforms and established publishers while disadvantaging smaller participants operating in lower-information contexts (Johnson et al., 2023). In equilibrium, advertisers trade off reach against attribution reliability, concentrating budgets in addressable environments and amplifying pre-existing disparities in monetization capacity (Aridor et al., 2023; P. Wang et al., 2024).

Distributional heterogeneity is equally pronounced on the user side. Privacy gains accrue primarily to users who exit tracking systems entirely, while privacy increasingly functions as a luxury good that is more readily secured by individuals with greater resources or capacity to navigate consent mechanisms (Dubé et al., 2025; Lin, 2022). By contrast, users who remain observable become more predictable within streamlined data environments and bear a greater share of surveillance and targeting intensity (Aridor et al., 2023; Zhao et al., 2023). Because consent design and behavioral constraints limit effective user control, privacy externalities are shifted toward those who do not, or cannot, exit. And reduced data availability may imply not only greater privacy for some users but also increased exclusion from data-driven allocation and service provision for others (Dubé et al., 2025).

Finally, redistribution under the GDPR extends beyond concentration to the opacity of data governance. As data collection retreats into closed, vertically integrated ecosystems, internal data

practices become less observable to regulators, competitors, and users (Peukert et al., 2022; Urban et al., 2020). While this shift reduces certain forms of widespread data leakage, it also weakens external scrutiny and amplifies concerns about accountability and systemic risk in highly centralized data environments (Johnson et al., 2023; Urban et al., 2020). Distributional power thus shifts not only toward larger actors, but toward less transparent modes of data governance (Peukert et al., 2022).

Taken together, GDPR enforcement increases average privacy while redistributing both privacy protection and economic surplus across the advertising value chain. The regulation curtails some forms of surveillance, strengthens data security, and reduces data intensity, but it also raises entry barriers, reallocates rents toward integrated and compliance-capable actors, and reshapes the long-run trajectory of data-driven markets. Privacy regulation therefore affects not only the level of data use, but the composition of participants, the allocation of surplus, and the structure through which data-driven economic activity evolves.

8 Limitations and Future Research

This thesis synthesizes a growing empirical literature, but several limitations should be noted. First, many empirical estimates are shaped by consent-based observability and regulatory spillovers. The presence of a recording wedge and the cross-border adoption of GDPR-style practices complicate causal interpretation and imply that reported effects often reflect bounds rather than precise point estimates of the regulation's true impact (Aridor et al., 2023; Congiu et al., 2022; Goldberg et al., 2024; Peukert et al., 2022).

Second, regulatory spillovers, often referred to as the "Brussels Effect", cause multinational firms to apply GDPR-compliant data practices globally to standardize infrastructure and reduce compliance complexity (Peukert et al., 2022). As a result, non-EU firms and websites may have voluntarily reduced third-party tracking following GDPR enforcement (Johnson et al., 2023), partially treating control groups in difference-in-differences designs. This contamination likely attenuates estimated treatment effects, implying that reported results represent conservative lower bounds of the GDPR's true impact (Lefrere et al., 2025; Miller et al., 2025).

Third, much of the available evidence captures short- to medium-run adjustments. While incumbent performance and content supply appear relatively stable, the long-run consequences for entry, innovation, and dynamic efficiency remain uncertain and may only materialize gradually over time (Janßen et al., 2022; Jia et al., 2021; Jones & Tonetti, 2020). In particular, current evidence is limited in its ability to assess how persistent data constraints affect experimentation, market contestability, and technological progress over longer horizons.

These limitations point to several priorities for future research. A central question concerns whether privacy-enhancing technologies, such as differential privacy, federated learning, or platform-led privacy infrastructures, can relax the trade-off between privacy and efficiency, or whether they instead reinforce advantages for large, vertically integrated platforms (Dubé et al., 2025; Johnson, 2024). More broadly, future work should examine how the design and implementation of privacy regulation shape dynamic competition, including entry, innovation, and market concentration, rather than focusing solely on static privacy or performance outcomes. Related research should also study the interaction between privacy regulation and competition policy, particularly in light of complementary interventions such as the Digital Markets Act (Johnson et al., 2023; Peukert et al., 2022).

9 Conclusion

This thesis examined how the General Data Protection Regulation reshaped privacy and economic performance in online advertising markets. Drawing on recent empirical evidence, it assessed whether the regulation achieved its stated privacy objectives and how these outcomes interacted with efficiency, firm behavior, and market structure.

The evidence shows that the GDPR increased privacy primarily by reducing data intensity and reshaping data governance rather than by establishing broad-based user control. Average privacy protection rose, but gains were unevenly distributed, reflecting heterogeneity in users' ability to exit tracking environments and firms' capacity to adapt to compliance requirements. Privacy gains were accompanied by real efficiency costs. Advertising effectiveness declined, and data-intensive entry and experimentation were constrained, raising concerns about dynamic efficiency and long-run innovation. These effects were not uniform: smaller, advertising-dependent firms were disproportionately affected, while large, vertically integrated platforms proved more resilient due to scale, first-party data, and compliance capacity.

Taken together, the findings highlight a fundamental trade-off. The GDPR successfully curtailed some forms of surveillance and reduced aggregate observability, but in doing so it reshaped the allocation of data access, surplus, and market power across the advertising ecosystem. Privacy regulation therefore influences not only the level of data use, but also the structure and long-run evolution of data-driven markets. Whether this constitutes a desirable outcome ultimately depends on normative judgments about the value of privacy relative to efficiency and innovation, as well as on firms' underlying revenue models and the extent to which privacy regulation expands participation rather than merely constraining data use. While the GDPR delivers meaningful privacy gains, these gains arise through compliance-driven retrenchment and market reorganization rather than sustained user empowerment. As a result, the regulation risks weakening competitive dynamics and innovation on dynamic margins, even as it strengthens privacy protection on average.

References

- Acar, G., Eubank, C., Englehardt, S., Juarez, M., Narayanan, A., & Diaz, C. (2014). The Web Never Forgets: Persistent Tracking Mechanisms in the Wild. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (pp. 674–689). Association for Computing Machinery. doi: 10.1145/2660267.2660347
- Acquisti, A., Taylor, C., & Wagman, L. (2016). The Economics of Privacy. *Journal of Economic Literature*, 54(2), 442–492. doi: 10.1257/jel.54.2.442
- Anderson, S. P., & Gabszewicz, J. J. (2006). The Media and Advertising: A Tale of Two-Sided Markets. In V. A. Ginsburg & D. Throsby (Eds.), *Handbook of the Economics of Art and Culture* (Vol. 1, pp. 567–614). Elsevier. doi: 10.1016/S1574-0676(06)01018-0
- Aridor, G., Che, Y.-K., & Salz, T. (2023). The effect of privacy regulation on the data industry: empirical evidence from GDPR. *The RAND Journal of Economics*, 54(4), 695–730. doi: 10.1111/1756-2171.12455
- Athey, S., Catalini, C., & Tucker, C. (2017). The Digital Privacy Paradox: Small Money, Small Costs, Small Talk (Tech. Rep. No. 23488). Cambridge, MA: National Bureau of Economic Research. doi: 10.3386/w23488
- Athey, S., & Gans, J. S. (2010). The Impact of Targeting Technology on Advertising Markets and Media Competition. *American Economic Review*, 100(2), 608–613.
- Bessen, J., Impink, S., Reichensperger, L., & Seamans, R. (2020). GDPR and the Importance of Data to AI Startups. *SSRN Electronic Journal*. doi: 10.2139/ssrn.3576714
- Campbell, J., Goldfarb, A., & Tucker, C. (2015). Privacy Regulation and Market Structure. *Journal of Economics & Management Strategy*, 24(1), 47–73. doi: 10.1111/jems.12079
- Chen, Z. (2022). Privacy Costs and Consumer Data Acquisition: An Economic Analysis of Data Privacy Regulation. *SSRN Electronic Journal*.
- Choe, C., Matsushima, N., & Shekhar, S. (2025). The Bright Side of the GDPR: Welfare-Improving Privacy Management. *Management Science*, 71(8), 6836–6858. doi: 10.1287/mnsc.2024.06653
- Choi, J. P., Jeon, D.-S., & Kim, B.-C. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics*, 173, 113–124. doi: 10.1016/j.jpubeco.2019.02.001.
- Congiu, R., Sabatino, L., & Sapi, G. (2022). The Impact of Privacy Regulation on Web Traffic: Evidence From the GDPR. *Information Economics and Policy*, 61, 101003. doi: 10.1016/j.infoecopol.2022.101003

- Das Chaudhury, R., & Choe, C. (2023). Digital Privacy: GDPR and Its Lessons for Australia. *Australian Economic Review*, 56(2), 204–220. doi: 10.1111/1467-8462.12506
- Degeling, M., Utz, C., Lentzsch, C., Hosseini, H., Schaub, F., & Holz, T. (2019). We Value Your Privacy ... Now Take Some Cookies: Measuring the GDPR's Impact on Web Privacy. In *Proceedings of the 2019 Symposium on Network and Distributed System Security (NDSS)*. doi: 10.14722/ndss.2019.23378
- Demirer, M., Jimenez Hernandez, D. J., Li, D., & Peng, S. (2024, February). Data, Privacy Laws, and Firm Production: Evidence from the GDPR (Working Paper No. 32146). National Bureau of Economic Research. doi: 10.3386/w32146
- Dubé, J.-P., Lynch, J. G., Bergemann, D., Demirer, M., Goldfarb, A., Johnson, G., . . . Tucker, C. (2025). Frontiers: The Intended and Unintended Consequences of Privacy Regulation for Consumer Marketing. *Marketing Science*, 44(5), 975–984. doi: 10.1287/mksc.2024.0901
- Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review*, 97(1), 242–259. doi: 10.1257/aer.97.1.242
- Eggers, F., Beke, F. T., Verhoef, P. C., & Wieringa, J. E. (2023). The Market for Privacy: Understanding How Consumers Trade Off Privacy Practices. *Journal of Interactive Marketing*, 58(4), 341–360. doi: 10.1177/10949968221140061
- European Parliament and Council of the European Union. (1995). Directive 95/46/ec. *Official Journal of the European Communities*. Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:31995L0046> (Directive 95/46/EC, OJ L 281, 23 November 1995)
- European Parliament and Council of the European Union. (2002). Directive 2002/58/ec. *Official Journal of the European Communities*. Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32002L0058> (Directive 2002/58/EC, OJ L 201, 31 July 2002)
- European Parliament and Council of the European Union. (2016). General Data Protection Regulation (GDPR). *Official Journal of the European Union*. Retrieved from <https://eur-lex.europa.eu/eli/reg/2016/679/oj> (Regulation (EU) 2016/679, OJ L 119, 4 May 2016)
- Farahat, A., & Bailey, M. C. (2012). How Effective Is Targeted Advertising? In *Proceedings of the 21st International Conference on World Wide Web* (pp. 111–120).
- Frey, C. B., & Presidente, G. (2024). Privacy Regulation and Firm Performance: Estimating the GDPR Effect Globally. *Economic Inquiry*, 62(3), 1074–1089. doi: 10.1111/ecin.13213

- Gal, M. S., & Aviv, O. (2020). The Competitive Effects of the GDPR. *Journal of Competition Law & Economics*, 16(3), 349–391. doi: 10.1093/joclec/nhaa012
- Godinho de Matos, M., & Adjerid, I. (2022). Consumer Consent and Firm Targeting After GDPR: The Case of a Large Telecom Provider. *Management Science*, 68(5), 3330–3378. doi: 10.1287/mnsc.2021.4054
- Goldberg, S. G., Johnson, G. A., & Shriver, S. K. (2024). Regulating Privacy Online: An Economic Evaluation of the GDPR. *American Economic Journal: Economic Policy*, 16(1), 325–358. doi: 10.1257/pol.20210309
- Goldfarb, A., & Tucker, C. E. (2011). Privacy Regulation and Online Advertising. *Management Science*, 57(1), 57–71. doi: 10.1287/mnsc.1100.1246
- Hagiu, A., & Wright, J. (2023). Data-enabled learning, network effects, and competitive advantage. *RAND Journal of Economics*, 54(4), 639–670. doi: 10.1111/1756-2171.12453
- Hils, M., Woods, D. W., & Böhme, R. (2020). Measuring the Emergence of Consent Management on the Web. In *Proceedings of the ACM Internet Measurement Conference* (pp. 317–332). doi: 10.1145/3419394.3423647
- Janßen, R., Kesler, R., Kummer, M. E., & Waldfogel, J. (2022, May). GDPR and the Lost Generation of Innovative Apps (Working Paper No. 30028). National Bureau of Economic Research. doi: 10.3386/w30028
- Jha, N., Trevisan, M., Mellia, M., Fernandez, D., & Irrarazaval, R. (2025). Privacy Policies and Consent Management Platforms: Growth and Users' Interactions over Time. *ACM Transactions on the Web*, 19(3). doi: 10.1145/3725737
- Jia, J., Jin, G. Z., & Wagman, L. (2021). The Short-Run Effects of the General Data Protection Regulation on Technology Venture Investment. *Marketing Science*, 40(4), 661–684. doi: 10.1287/mksc.2020.1271
- Johnson, G. A. (2024). Economic Research on Privacy Regulation: Lessons from the GDPR and Beyond. In A. Goldfarb & C. E. Tucker (Eds.), *The Economics of Privacy* (pp. 97–126). Chicago: University of Chicago Press. doi: 10.7208/chicago/9780226834085-008
- Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer Privacy Choice in Online Advertising: Who Optes Out and at What Cost to Industry? *Marketing Science*, 39(1), 33–51. doi: 10.1287/mksc.2019.1198
- Johnson, G. A., Shriver, S. K., & Goldberg, S. G. (2023). Privacy and Market Concentration: Intended and Unintended Consequences of the GDPR. *Management Science*, 69(10), 5695–5721. doi: 10.1287/mnsc.2023.4709

- Jones, C. I., & Tonetti, C. (2020). Nonrivalry and the Economics of Data. *American Economic Review*, 110(9), 2819–2858. doi: 10.1257/aer.20191330
- Ke, T. T., & Sudhir, K. (2023). Privacy Rights and Data Security: GDPR and Personal Data Markets. *Management Science*, 69(8), 4389–4412. doi: 10.1287/mnsc.2022.4614
- Lefrere, V., Warberg, L., Cheyre, C., Marotta, V., & Acquisti, A. (2025). Does Privacy Regulation Harm Content Providers? A Longitudinal Analysis of the Impact of the GDPR. *Management Science*. doi: 10.1287/mnsc.2022.03186
- Libert, T., Graves, L., & Nielsen, R. K. (2018). Changes in Third-Party Content on European News Websites after GDPR (RISJ Factsheet). Reuters Institute for the Study of Journalism. doi: 10.60625/risj-r2ex-7248
- Lin, T. (2022). Valuing Intrinsic and Instrumental Preferences for Privacy. *Marketing Science*, 41(4), 663–681. doi: 10.1287/mksc.2022.1368
- Lin, T., & Strulov-Shlain, A. (2025). Choice Architecture, Privacy Valuations, and Selection Bias in Consumer Data. *Marketing Science*, 44(6), 1321–1338. doi: 10.1287/mksc.2023.0373
- Lu, X., Zhao, X., & Xue, L. (2016). Is Combining Contextual and Behavioral Targeting Strategies Effective in Online Advertising? *ACM Transactions on Management Information Systems*, 7(1), 1–20. doi: 10.1145/2883816
- Miller, K. M., Lukic, K., & Skiera, B. (2025). The Impact of the General Data Protection Regulation (GDPR) on Online Tracking. *International Journal of Research in Marketing*. (In press, corrected proof. Available online 11 March 2025) doi: 10.1016/j.ijresmar.2025.03.002
- Neumann, N., Tucker, C. E., Kaplan, L., Mislove, A., & Sapiezynski, P. (2024). Data Deserts and Black Boxes: The Impact of Socio-Economic Status on Consumer Profiling. *Management Science*, 70(11), 8003–8029. doi: 10.1287/mnsc.2023.4979
- Nissenbaum, H. (2010). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press.
- Nokhbeh Zaeem, R., & Barber, K. S. (2020). The Effect of the GDPR on Privacy Policies: Recent Progress and Future Promise. *ACM Transactions on Management Information Systems*, 12(1), 1–20. doi: 10.1145/3389685
- Nouwens, M., Liccardi, I., Veale, M., Karger, D., & Kagal, L. (2020). Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery. doi: 10.1145/3313831.3376321

- Peukert, C., Bechtold, S., Batikas, M., & Kretschmer, T. (2022). Regulatory Spillovers and Data Governance: Evidence from the GDPR. *Marketing Science*, 41(4), 746–768. doi: 10.1287/mksc.2021.1339
- Rafieian, O., & Yoganarasimhan, H. (2021). Targeting and Privacy in Mobile Advertising. *Marketing Science*, 40(2), 193–218. doi: 10.1287/mksc.2020.1235
- Smith, M., Torres-Agüero, A., Grossman, R., Sen, P., Chen, Y., & Borcea, C. (2024). A Study of GDPR Compliance under the Transparency and Consent Framework. In *Proceedings of the ACM Web Conference 2024* (pp. 1227–1236). doi: 10.1145/3589334.3645618
- Sun, T., Yuan, Z., Li, C., Zhang, K., & Xu, J. (2024). The Value of Personal Data in Internet Commerce: A High-Stakes Field Experiment on Data Regulation Policy. *Management Science*, 70(4), 2645–2660. doi: 10.1287/mnsc.2023.4828
- Tucker, C. (2023). Algorithmic Exclusion: The Fragility of Algorithms to Sparse and Missing Data (Tech. Rep.). Washington, DC: Brookings Institution, Center on Regulation and Markets.
- Urban, T., Tatang, D., Degeling, M., Holz, T., & Pohlmann, N. (2020). Measuring the Impact of the GDPR on Data Sharing in Ad Networks. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security (Asia CCS)* (pp. 222–235). Association for Computing Machinery. doi: 10.1145/3320269.3372194
- Utz, C., Degeling, M., Fahl, S., Schaub, F., & Holz, T. (2019). (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (pp. 973–990). Association for Computing Machinery. doi: 10.1145/3319535.3354212
- Varian, H. R. (2007). Position Auctions. *International Journal of Industrial Organization*, 25(6), 1163–1178. doi: 10.1016/j.ijindorg.2006.10.002
- Veale, M., & Zuiderveen Borgesius, F. (2022). Adtech and Real-Time Bidding under European Data Protection Law. *German Law Journal*, 23(2), 226–256. doi: 10.1017/glj.2022.18
- W3C Tracking Protection Working Group. (2019). Tracking compliance and scope (Working Group Note). World Wide Web Consortium. Retrieved from <https://www.w3.org/TR/2019/NOTE-tracking-compliance-20190122/>
- Wang, J., Zhang, W., & Yuan, S. (2017). Display Advertising with Real-Time Bidding (RTB) and Behavioural Targeting (Tech. Rep. No. 1610.03013). arXiv. Retrieved from <https://arxiv.org/abs/1610.03013>
- Wang, P., Jiang, L., & Yang, J. (2024). The Early Impact of GDPR Compliance on Display Advertising: The Case of an Ad Publisher. *Journal of Marketing Research*, 61(1), 70–91. doi: 10.1177/00222437231171848

- Westin, A. F. (1967). *Privacy and freedom*. Atheneum.
- Yan, J., Liu, N., Wang, G., Zhang, W., Jiang, Y., & Chen, Z. (2009). How Much Can Behavioral Targeting Help Online Advertising? In *Proceedings of the 18th International Conference on World Wide Web* (pp. 261–270).
- Yuan, S., Wang, J., & Zhao, X. (2013). Real-Time Bidding for Online Advertising: Measurement and Analysis. arXiv preprint. Retrieved from <https://arxiv.org/abs/1306.6542>
- Zhang, K., & Katona, Z. (2012). Contextual Advertising. *Marketing Science*, 31(6), 980–994. doi: 10.1287/mksc.1120.0740
- Zhao, Y., Yildirim, P., & Chintagunta, P. K. (2023). Privacy Regulations and Online Search Friction: Evidence from GDPR (Tech. Rep. No. 23-141). Marketing Science Institute.

Appendices

Appendix 1 Explanation of the use of AI

The following table documents how generative AI tools were used in the preparation of this thesis, in accordance with the University of Turku guideline on artificial intelligence in teaching and studying available at <https://utuguides.fi/artificialintelligence>. The tools were used in a way that complies with academic integrity guidelines; all AI output was reviewed, verified, and edited by the author, who takes full responsibility for the content of the thesis.

Chapter	AI tool	Version	Prompt	Explanation / Comment
All	NotebookLM	Not specified	List the key papers related to heterogeneous privacy effects and data externalities in digital markets, and group them by topic.	Used to identify and organize relevant academic literature that I uploaded to NotebookLM.
All	ChatGPT	5.2	Reword the following paragraph in a more academic tone, without changing meaning or adding new content: In online advertising, these concerns are related to how much firms can actually observe user behavior...	In online advertising, these normative concerns map naturally to the extent to which firms can observe user behavior...
All	ChatGPT	5.2	Please analyze the following section and flag any logical inconsistencies, unclear formulations, or grammar errors and return it.	Returned text was reviewed and selectively revised based on the tool's feedback.
4.1	ChatGPT	5.2	Based on the following description, suggest some appropriate academic terms or labels for this mechanism: After the GDPR was introduced, firms responded by mainly....	Looked at output and decided for one of the labels: 1. Risk-Based Retrenchment. Why it fits... 2. Selective data retrenchment... 3. Compliance-driven data pruning...