A Systematic Review of User Experiments on the Effects of Dark Patterns

Brennan Schaffner
Georgetown University
Washington, District of Columbia
USA
brennan.schaffner@georgetown.edu

Luis Heysen
Department of Computer Science
University of Chicago
Chicago, Illinois, USA
lheysent@uchicago.edu

Marshini Chetty
Department of Computer Science
University of Chicago
Chicago, Illinois, USA
marshini@uchicago.edu

Abstract

Deceptive/Manipulative Patterns (DMP) are interface designs, also known as "dark patterns," that manipulate user behavior. While considerable attention has been paid to their ethical and legal implications, empirical evidence about their real-world effects remains diffuse. This review synthesizes up-to-date experimental studies, focusing on works that quantify how (or whether) DMPs influence users. We also aggregate findings on interventions aimed at reducing DMP effects. Our synthesis highlights the experimental agreement that DMPs do significantly alter user behavior (with large variance in effect size) and that external interventions have been mostly unsuccessful in mitigating their effects. Lastly, we show that significant correlations between DMP effects and personal characteristics (e.g., age or political affiliation) are uncommon, indicating DMPs similarly affected nearly all populations tested. By summarizing the experimental evidence, we clarify the effects of DMPs, highlight gaps and tensions in the existing experimental literature, and help inform ongoing research and policy directions.

CCS Concepts

• **Human-centered computing → Empirical studies in HCI**; • **Security and privacy → Human and societal aspects of security and privacy**.

Keywords

manipulative design, dark patterns, deceptive design, deceptive/manipulative pattern, user experiments, systematic review

ACM Reference Format:
Brennan Schaffner, Luis Heysen, and Marshini Chetty. 2026. A Systematic Review of User Experiments on the Effects of Dark Patterns. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3772318.3790383

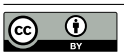

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3790383

1 Introduction

Deceptive/Manipulative Patterns (DMPs)[1] are "user interface design choices that benefit an online service by coercing, steering, or deceiving users into making decisions that, if fully informed and capable of selecting alternatives, they might not make" [47]. Common examples include platforms employing excessively difficult processes to cancel online subscriptions or change privacy-unfriendly default settings. DMPs have provoked mounting scrutiny from scholars, regulators, and practitioners, and evidence about their measurable impact on users steadily accumulates [13, 22, 53, 54]. There is general agreement about DMPs' pervasiveness and harmful effects. This sentiment is supported by several important studies, such as the unifying of foundational taxonomies into a cohesive ontology [27], measuring the ubiquity of DMPs [17, 47], and studying their potential harms [48]. Yet, skepticism duly remains. It is not well understood whether DMPs have quantified user harms. Accordingly, from the perspective of industry and marketing practitioners, DMPs are often considered standard practice, reflecting status quo sales strategies that have long been prevalent in commerce [5, 48].

The researcher response has been to conduct properly controlled experimentation on the effects of DMPs. There are many challenges to conducting studies that isolate the effects of DMPs threatening ecological validity, including operationalizing DMPs in near real-world environments while considering potential participant harms. Nevertheless, researchers have designed studies to successfully isolate the effects of DMPs. For example, an early publication in DMP scholarship conducted a randomized control trial experiment quantifying the effect that different DMPs had on the rates in which consumers stayed enrolled in a malicious subscription [42].

Official investigations and court actions have concluded that leaving DMPs unchecked can result in consumer harm. The United States Federal Trade Commission's 2022 report on DMPs describes numerous accounts of manipulative designs extracting millions of dollars from users and tricking them into sharing personal data [23]. In response, there has been a clear rise in regulatory actions against DMPs [20, 23, 56], sometimes manifesting in high-profile cases. For instance, Epic Games agreed to pay a $245 million dollar settlement for using DMPs to trick children into making unwanted purchases [24]. The accelerating regulatory landscape highlights the role of experimental evidence informing policy. Although such experiments have entered the landscape, the results are scattered across disciplines and venues, leaving regulators and researchers

[1] We use Deceptive/Manipulative Patterns (DMPs) to refer to "dark patterns" as currently recommended by the Association for Computing Machinery (ACM) [18]. For the sake of discoverability and connection to prior work, we retain "dark pattern" in the title and as a keyword given its common use by researchers, legislators, and litigators.

without a clear, cumulative picture of what is known and where the gaps lie.

The present work quantitatively assists the discussions around whether DMPs change user behavior by synthesizing relevant, to-date experimental evidence supporting (and, to a minor extent, contesting) the harmful effects that DMPs have on users. Previous reviews have targeted adjacent yet distinct topics, broadly reviewing designs that influence behavior [59, 64]—especially those aimed to *help* users rather than harm them [2, 34]—or small subsets of DMPs specifically [9, 32, 51, 67]. By contrast, our review captures the breadth of all current scholarship that experimentally measures DMP harms.

Building on this synthesis, our review also evaluates the state of DMPs experimental work with four additional analytical angles: We (i) synthesize results from experiments that have tested potential interventions aimed at combating the effects of DMPs, such as DMP educational sessions or interstitials that promote reflection, (ii) aggregate experiments that have tested marginal effects of "stacking" multiple DMPs, (iii) examine which types of DMPs have been tested more frequently and whether specific instances are more effective than others, and (iv) summarize evidence on whether users' personal characteristics mediate the effects of DMPs.

In sum, our review of experimental DMP literature holds the following objectives:

- In which areas do experimental results agree and/or diverge with respect to quantitative measures of DMPs' effects as well as interventions designed to counteract them?
- What aspects of DMPs have been well described with quantitative measures, and what aspects have been understudied?

Considering the increased regulatory efforts aimed at curbing the use of DMPs,[2] answering these questions is important for aligning the DMPs community and guiding legislative directions. Our review finds experimental agreement that DMPs commonly have measurable effects on user behavior and that interventions have been thus far ineffective mitigators. We also reveal which types of DMPs have been studied most thoroughly and which have been understudied experimentally. Ultimately, our review helps researchers and policymakers focus future research and regulatory efforts where they are most needed.

2 Related Work

We summarize the present field of DMPs and the relevant previously conducted systematic reviews.

2.1 Research On DMPs

Deceptive and manipulative digital interfaces have become a popular matter of focus across regulatory [13, 33, 56], academic [29, 53], and design spaces [22, 60]. Recent legislation has made efforts to specifically define and ban DMPs in specific contexts. For example, two of the earliest legal codifications of DMPs are in the California Privacy Rights Act in the United States (U.S.) which defines DMPs as "user interface[s] designed or manipulated with the substantial effect of subverting or impairing user autonomy, decision making, or choice" [63], and the European Union's Digital Services Act which define DMPs as "[online] practices that materially distort or impair, either on purpose or in effect, the ability of recipients of the service to make autonomous and informed choices" [21].

Scholars have substantially elaborated upon the harms that DMPs may cause, including harms to individual welfare, collective welfare, and individual autonomy [48]. For example, cookie consent banners may bury privacy-friendly options and employ confusing controls resulting in users "opting" for more data to be shared with the platform and third-parties.

Taxonomic efforts have iterated and evolved the classifications of DMPs since they were originally cataloged in 2010 [14]. Gray et al.'s contemporary ontology published in 2024 consists of five High-Level DMP strategies: Social Engineering, Obstruction, Sneaking, Interface Interference, and Forced Action [27]. Despite their potential for harm, DMPs have been found to be quite prevalent. In 2020, researchers found DMPs in 95% of 240 popular mobile apps from the Google Play Store [17] and in 89% of cookie banners in the EU [54].

To measure the effects that DMPs have on users, researchers have begun employing controlled experiments, the earliest of which—published in 2019—studied how 80,000 users of a German website interacted with its consent banner [66]. They found that cookie banner implementation using DMPs resulted in significantly higher consent rates than those without these patterns. Experiments have also demonstrated the potential for financial harms from DMPs [38, 42, 70]. For example, an early DMP experiment from the field of law revealed the power of DMPs in tricking consumers into enrolling in a dubious paid subscription [42]. Our work is the first systematic review aggregating the results of all such experiments—those that measure the effects of DMPs—to date.

2.2 Relevant Extant Systematic Reviews of DMPs

Given the growth in scholarship around DMPs, there has been a corresponding effort to conduct systematic literature reviews to summarize findings and identify new research angles. However, no prior systematic review has analyzed the results of all peer-reviewed, experimental studies across the domain of DMPs. In this section, we discuss the systematic reviews that this work builds upon and highlight the need for the present work.

Reviews that Broadly Address DMPs. Multiple works have reviewed and synthesized research on DMPs, aiming to define, classify, and characterize how manipulative design techniques are studied and understood in digital platforms. In a work-in-progress paper, Gray et al. conducted a general review outlining the high-level components of DMPs literature [29]. They found that the most common DMP contributions took the form of *describing* or *framing* DMPs, and the most common methodologies used were content analysis, followed by experiments (for which they identified 14 works at the time). Yi and Li conducted a DMP systematic review focusing on the wide array of regulatory approaches, advocating for a paradigm shift from harm-based frameworks to proactive alternatives [69]. Several systematic reviews focused on synthesizing identified DMPs into a cohesive, standard taxonomy [27, 48, 58]. These existing taxonomies informed our approach: specifically, in our methods (see §4.2), we use Gray et al.'s ontology to categorize

[2] https://www.deceptive.design/cases

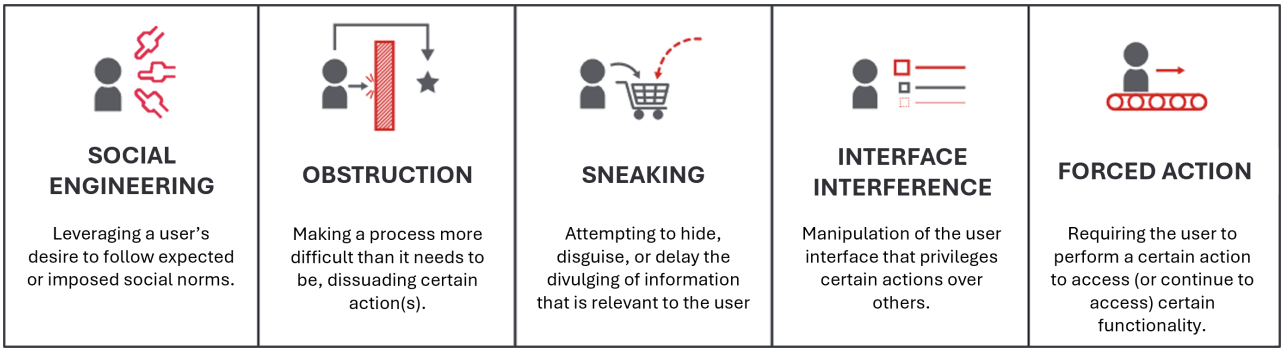

Figure 1: The five high-level categories in Gray et al.'s ontology. Each high-level category also comprises meso- and low-level patterns. This is an edited figure from Gray et al. [28] updated to match the contemporary ontology [27].

DMPs present in the experimental papers reviewed. See Figure 1 for abridged definitions of the five high-level DMP categories.

Independent and government organizations have also released technical reports including substantial reviews that outline many findings relevant to DMPs (e.g., [56]). These reports tend to cover a wide range of information and apply less defined criteria for determining which works to include, resulting in reports that are broadly educational and less reproducible compared to the systematic review of experiments on DMPs conducted in the present work.

Reviews of Influential Design. There are also a number of related systematic reviews that do not rely on the framing of DMPs but have peripheral overlap in contribution to the DMPs space. For instance, there are several reviews focusing on the concept of "nudging" [2, 34, 35]. Acquisti et al. review privacy and security literature related to nudging looking broadly at the limitations and potential of existing interventions and give guidelines for ethical nudge design [2]. Multiple studies have been conducted that quantitatively reviewed nudging literature to assess the effectiveness of nudging interventions in general and in different application contexts [34, 35]. Both DMPs and nudging could be considered theoretical subsets of persuasive technology [59] or automated influence [64], both of which have been systematically reviewed for various ethical considerations [59, 64].

In our review, we focus specifically on DMPs literature, which we distinguish from nudging following previous literature (e.g., [8]), where nudges are "soft paternalism" that guide users towards better decisions [1, 2] and dark pattens harm a user's ability to make informed decisions [27, 28, 48]. In short, while both are operationalized by changes to choice architecture, nudges are said to help users, whereas DMPs are said to harm users.

Context-Specific Reviews On DMPs. Some reviews have focused on specific subsets or contexts of DMPs. For instance, Bielova conducted a review of DMP effects on users' acceptance rates in cookie banners with the French Data Protection Authority [9]. Hadan et al. reviewed the DMP literature in specifically extended reality (XR) environments identifying risks and harms unique to XR [32]. Westin and Chiasson review DMPs literature through a paradigm of social pressures [67]. In another review, Roffarello et al. develop a framework for a specific subset of DMPs related to attentional harms, defining "attention capture damaging patterns" [51]. Our review captures multiple domains of study and instead focuses on experiments quantifying the effects of various DMPs.

Reviews On DMP Harms. Some reviews focus on the harms of DMPs [3, 48, 61]. For instance, Mathur et al. review DMPs literature and outline normative perspectives through which the harms of DMPs can be analyzed [48]. Sanju and Kumar align the types of DMPs with specifically how they harm user autonomy along four dimensions: agency, freedom of choice, control, and independence [3]. In another review, Cara sorted DMPs by their harm severity, from "just annoying" to "need official regulation" [15]. Most recently, Santos et al. review relevant DMPs literature to create a taxonomy of the possible harms that DMPs can cause and discuss the challenges surrounding the more complicated harms such as those that are non-material or societal [61]. In our review, we build on this literature to systematically look at studies that have *experimentally measured* these harms.

3 Review Scope

Here we outline the scope of our systematic review. We outline our inclusion parameters for the studies we ultimately analyzed as part of our dataset.

DMP Framing. The study must use the framing of DMPs, meaning the paper would have to contain the phrase "DMP"—or its parallel terms (i.e., "deceptive and manipulative design" and "deceptive design patterns"—in the work's framing. Effectively, this means the terminology had to appear in the title, abstract, keywords, or background sections to be included in our dataset.

Peer-Reviewed. We excluded studies that were not peer-reviewed, including technical reports, academic theses, and white papers. However, we included preprints that had gone through the peer-review process and were yet to appear at a pending conference. The inclusion of government reports (as discussed in the §4.1) is the only exception to this criterion or any other criterion in this section.

User Experience (UX) Elements That Lead To Harm. To distinguish from literature that uses persuasive design largely for the user's

benefit, we excluded studies that do not engage with known user harms. In practice, this exclusion separates the work most notably from those of "Nudge Theory," which is primarily housed in the behavioral sciences and originated with the intention of paternal aid for user decision making, especially according to those who popularized it [65]. Some researchers consider DMPs as an (evil) type of nudge. We do not exclude such studies since they retain the framing and harmful character of DMPs. This criterion also led to the exclusion of papers that measure "neutral" persuasive design, where the platform benefit and user harms are unclear (e.g., [19]).

Direct Measures of Behavior. In spirit of building the strongest suite of evidence, we excluded works that rely on self-reported metrics instead of direct measures of participant behavior. In practice, this marks the difference between what participants *actually do* when faced with DMPs and what they *say they would do* when presented with an example of a DMP in a survey. This also led to the exclusion of studies that used survey constructs or scales as experimental dependent variables (e.g., perceived platform trustworthiness) because they are indirect proxies for behavior rather than direct behavioral measures.

Dependent Variables Directly Relevant to DMPs. We excluded studies that examined the effects of DMPs on behavior that are tangential to the "intent" of the DMP. For example, some studies measure how DMPs increase users' "consent" to the sale of their personal data (a direct effect) as well as how the DMPs may affect how long it takes users to make the consent decision (an indirect effect). Studies about the indirect effects are excluded for the sake of simplifying the interpretation of the findings.

Statistical Significance on Quantitative Metrics. Due to industry responses to both enacted regulation and threats thereof that DMPs are not well-studied or definitively harmful to consumers, we chose to seek what is most popularly accepted to be "hard" evidence. That is, studies must test for statistical significance as well as report corresponding statistical measures (e.g., p-values) in order to be included in the analysis.

Unfortunately, these criteria led to the exclusion of papers that quantitatively or descriptively measured how DMPs affected behavior—often to a convincing extent—but did not calculate statistical significance. For example, Monge Roffarello and De Russis qualitatively found that removing certain DMPs on social media led to reductions in time spent on the platform [50].

The dataset ultimately included the scientific methodologies typically deemed most rigorous (i.e., "true experiments"), including between-subjects post-test designs with control groups, within-subjects repeated measures designs with counterbalancing, or a mix of both.

Such controlled and statistically-minded behavioral experiments included in the resulting dataset are not always practical, necessary, or superior in demonstrating the effects of deceptive and manipulative design [68]. Our choice to steel man the position of a DMP skeptic should not be interpreted as a value or relevancy judgment on the plethora of impressive and important studies that measure the effects of DMPs qualitatively or through expert or legal analysis. See §6.3 for further discussion on this matter.

Lastly, many works use a combination of included and excluded studies together in any single publication. We include only the relevant components from each work. When multiple works were found to include duplicately-published content, relevant results were only included once.

4 Methodology

Here we present the processes employed to collect the corpus of experiments fitting our review criteria and how they were analyzed.

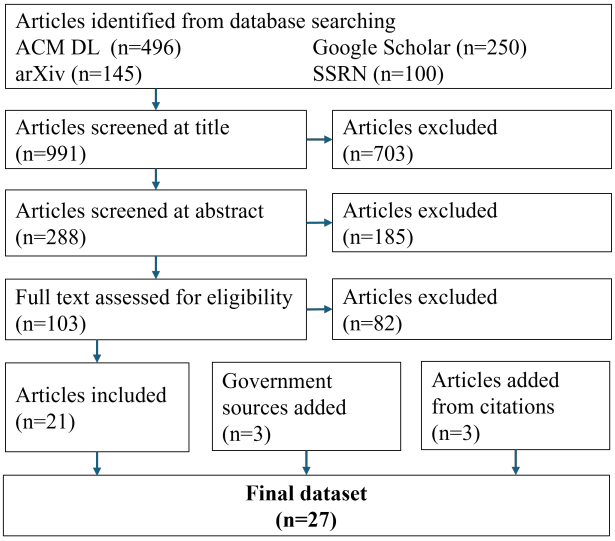

Figure 2: Flow of study selection through the screening process, from 991 initial papers filtered down to the final 27.

4.1 Corpus Creation

We assembled a comprehensive corpus of experimental literature on DMPs by systematically searching multiple academic repositories between January and May 2025.[3] Our repository selection ensured coverage across technical, legal, and social science fields.

For breadth across disciplines, we employed a Python-based Selenium web scraper to retrieve the first 250 results from **Google Scholar**, in line with prior work [29].[4] For computer science publications, our search returned 496 results in the **ACM Digital Library**,[5] which we retrieved using the platform's native batch download feature. To capture works in the adjacent fields of law and social science research, we downloaded the first 100 results from **SSRN**,[6] reflecting its more specialized scope and tangential relevance to Human-Computer Interaction. Lastly, we obtained the 145 publications our query returned in **arXiv**[7] to ensure technical coverage and capture to-be-published preprints.

All search results were exported as BibTeX files and consolidated using Rayyan.ai,[8] a collaborative tool for systematic reviews and

[3] Search queries are provided in Appendix A.1.
[4] https://scholar.google.com/
[5] https://dl.acm.org/
[6] https://www.ssrn.com/index.cfm/en/
[7] https://arxiv.org/
[8] https://rayyan.ai/

Table 1: The publication domain, venue, and year for the 27 final papers in the dataset. The domain column also shows the distribution of experimental units across domains in parentheses.

Domain		Venue		Year	
Consent Popups	16 (51)	Conference	12	2025	4
Subscriptions	4 (41)	Journal	8	2024	4
Privacy Settings	2 (27)	Non-Archival	3	2023	6
E-Commerce	2 (24)	Government Study	3	2022	4
Streaming	1 (3)	Preprint	1	2021	6
Advertising	1 (1)			2020	2
Donations	1 (1)			2019	1

automated detection of duplicate papers. In total, our initial dataset contained 991 papers before screening.

4.1.1 Screening Process. We applied a structured, multi-stage screening process to identify only the relevant papers. At every stage of screening, we favored inclusion in cases of uncertainty. The process, summarized in Figure 2, consisted of the following steps:

(1) **Title/Abstract Screening**: From the initial dataset of 991 papers, titles and abstracts were reviewed, and papers clearly not relevant to DMPs or experimental approaches were excluded, resulting in 103 papers. Duplicate papers were also excluded. The most common reasons for exclusion at this stage were irrelevance arising from overlap in search terms with popular concepts in other fields (e.g., 'dark matter').
(2) **Full-Text Screening**: Full texts were assessed according to the inclusion and exclusion criteria specified in §3. Three papers had to be retrieved directly from the authors but were ultimately excluded for not meeting the scope criteria. Papers unavailable in English were also excluded. The most common reasons for exclusion at this stage were: not conducting controlled experiments, not being peer-reviewed, or measuring user perceptions rather than behavior change. An accurate count of exclusion reasons is unavailable since analysis was stopped upon encountering the first instance of disqualifying criteria, so not every disqualifying characteristic was recorded.
(3) **Government Reports**: Based on domain expertise and expert consultation, three relevant government studies meeting the inclusion criteria were incorporated ([11, 16, 43]).[9]
(4) **Citation Search**: Reference lists of included papers were analyzed, with identified studies subjected to the same screening process, resulting in three additional papers.

A total of 27 papers were included in the final dataset.

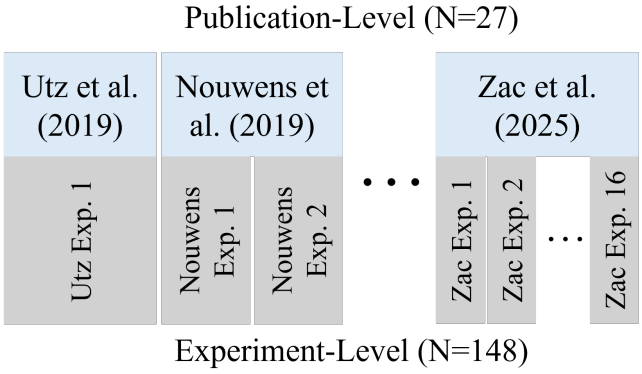

Figure 3: Visual depiction of both the publication- and experiment-levels in the dataset.

4.2 Corpus Analysis

We first recorded relevant meta details for each of the publications in the final dataset. For each publication in the corpus, we noted the domain (e.g., e-shopping or consent pop-ups), the venue type (e.g., peer-reviewed journal), and publication year, which can be seen in Table 1. Since each publication could include more than one experiment, we then developed a codebook to apply at the *experiment*-level as discussed below.

4.2.1 Identifying Experimental Units. To enable consistency across the dataset, we broke down multi-part studies into relevant *experimental units* (illustrated in Figure 3), where each *experimental unit* represents a full study fitting the inclusion criteria outlined in §3. For example, Graßl et al. [26] test the difference between several consent dialogues with permuted design against a neutral control dialogue. They test for statistically significant effects between participants in three DMPs conditions and a control condition. Hence, for the sake of our review, Graßl et al.'s work is reduced to three experimental units, one for each DMP effect reported. They also tested for effects of various conditions that "nudge towards privacy-friendly options" as well as how various conditions—both DMPs and nudging—affect participants' self reported levels of control and deliberation, all of which are out of scope of this review.

For each experimental unit, we noted the following:

- **DMP(s) Tested.** We used the leading ontology of DMPs from Gray et al. [27][10] to standardize the terminology across the corpus. In some cases, this required reclassifying the DMPs reported by the original authors into standard terminology. Moreover, if the authors were testing the effects of a design that did not fit into the ontology, it was excluded for not having scholarly consensus on being a DMP.
- **Experimental Results.** We coded whether the experiment resulted in statistically significant results, to what degree of significance, and the effect size if reported in universal

[9] Duplicate findings present in both Bogliacino and EU Commission (2022) are counted only once, attributed to the latter.

[10] Gray et al.'s taxonomy aggregates terminology from the ten most frequently cited academic and regulatory reports on DMPs.

Table 2: Summary statistics for the final dataset of 148 experimental units across 27 papers, including the recruitment methods used, the number of experimental units per paper, and the number of participants per experimental unit. Studies could use multiple recruitment methods. Note that the sample sizes reported are skewed higher since many papers reported only a total sample size without reporting the relevant subset sample sizes of each experimental conditions. Posner et al. conducted a large-scale natural experiment; therefore, sample size metrics are reported both including and excluding it (the latter shown in parentheses).

Recruiting Method	Exp. units / No. Papers	Exp. units Per Paper		Sample Size Per Exp. (without Posner et al.)	
Online recruiting firms	115 / 15	Median	3	Median	925 (922)
University connections or mailing lists	9 / 4	Average	5.48	Average	19,314 (1758)
		SD	6.46	SD	212,947 (6432)
Live A/B tests	9 / 4	Range	1 – 26	Range	40—2.6M (40–70,208)
Not specified	9 / 2				
Google ads	9 / 2				
Other	3 / 2				

terms. It is important to note that lacking statistical significance (failing to reject the null hypothesis) is not identical to proving no effect exists (affirming the null hypothesis).

- **Participant and Recruitment Details.** We recorded the sample size of each experimental unit, recruitment methods, devices used, whether specific demographics were targeted, and study locations.
- **Other Experimental Details.** We recorded the independent variables (e.g., presence of a DMP), dependent variable (e.g., data shared or product purchased), experiment type (e.g., posttest-only control group design), and experimental setting (e.g., in the wild or online study with deception).

All experimental units fell into one of the following five categories which are discussed separately in the following sections:

Category 1: tested for effects induced by presence of DMP(s) compared to a control without the DMP(s).

Category 2: tested whether external interventions mitigated the effects induced by DMPs, where external interventions are attempts to reduce DMP effects without removing them altogether.[11]

Category 3: tested for additive effects when multiple DMPs were present.

Category 4: compared the effects from distinct types of DMPs.

Category 5: tested for correlations between the effects of DMPs and user personal characteristics, ranging from demographics (e.g., age and gender) to personal attributes (e.g., technology affinity and political affiliation). These tests were commonly post-hoc analyses at a paper-level (rather than the experiment-level), and therefore these results are excluded from any experiment unit counts in the following sections. We instead present a summarization at the paper-level in §5.6.

[11]The removal of a DMP altogether constitutes a control group in a Category 1 experiment. What distinguishes external interventions is that they represent external systems or changes to the interface that are not directly changing the DMPs. Similarly, "bright patterns" are not considered external interventions since they operate within the same vein as the DMPs, essentially removing them altogether.

4.3 Methodological Limitations

Our systematic review methodology has several limitations that should be acknowledged. First, as a snapshot of literature collected up to May 2025, we cannot guarantee inclusion of papers published after our search period. Second, despite our comprehensive search strategy across multiple repositories, we may have missed rare papers not indexed in our selected databases or those using terminology outside our search parameters. We mitigated this limitation by consulting the cited works of the papers resulting from our search. Third, our review may be affected by publication bias, as studies with null results are less likely to be submitted and published, a well-known metascientific phenomenon [36]. Fourth, the chosen representation of experimental units treats experiments on equal grounds despite their individual differences in important parameters such as sample size. However, this bias is moderately counteracted by larger studies typically conducting more experiments and are thus weighted higher in analytical representation. Fifth, our search strategy employed cutoff points when reviewing ranked search results for two of the databases (Google Scholar and SSRN), introducing potential for missed relevant works. This risk was mitigated through conducting searches across multiple databases, setting search result thresholds beyond drop-offs in search result relevance, and consulting the citations of the included works for additional relevant studies to include. We took a conservative approach to minimize the potential of missing relevant works as evidenced by the high exclusion rate across screening stages. Last, the inclusion of only English-language publications may have excluded relevant work from non-English speaking regions, potentially limiting the diversity of perspectives in our review.

5 Findings

The dataset ultimately consisted of 148 experimental units from 27 papers, published from 2019 to 2025. A summary of the dataset is reported in Table 2, which shows the most common recruiting methods, distribution of experimental units per paper, and the

Table 3: Breakdown of the the different types of DMPs tested across the corpus at the experiment-level using Gray et al.'s ontology. Experimental units often implemented more than one DMP, hence the column totals surpassing the total experiment count. The numbers in this table represent the total unique instances of each DMP tested in the dataset.

High-Level		Meso-Level		Low-Level	
Interface Interference	157				
		Manipulating Choice Architecture	107	Visual Prominence	61
				False Hierarchy	46
		Bad Defaults	28		
		Emotional or Sensory Manipulation	13	Positive or Negative Framing	13
		Trick Questions	9		
Social Engineering	76				
		Shaming	28	Confirmshaming	28
		Urgency	19	Countdown Timer	9
				Limited Time Message	5
				Activity Message	5
		Personalization	11		
		Scarcity and Popularity Claims	9	High Demand	9
		Social Proof	9	Low Stock	9
Obstruction	45				
		Adding Steps	45	-	40
				Privacy Maze	5
Forced Action	21	-	1		
		Nagging	15		
		Attention Capture	3	Autoplay	3
		Forced Communication or Disclosure	2		
Sneaking	10				
		Hiding Information	9		
		Bait and Switch	1	Disguised Ads	1

distribution in number of participants included for each experiment. See Appendix A.2 Table 7 for a full list of the 27 papers that made the final cut and how many experimental units are included from each paper.

5.1 DMP Experiments Have Covered All High-Level DMP Types

In total, the 148 experiments tested the effects for a range of DMP types. Table 3 shows the breakdown of the DMPs in the dataset of experiments, at the high-, meso-, and low-levels [27]. All five high-level patterns appear in the dataset but to uneven degrees. Interface Interference was the most commonly tested high-level pattern, showing up 157 times in the dataset ($N = 85/148$ exp. units; 22/27 papers)—most of which are Visual Prominence or False Hierarchy. 76 instances of Social Engineering were tested ($N = 45/148$ e.; 6/27 p.), followed by 45 instances of Obstruction ($N = 41/148$ e.; 10/27 p.), and 21 instances of Forced Action ($N = 21/148$ e.; 6/22 p.). Only 10 instances of Sneaking were tested in the dataset ($N = 10/148$ e.; 4/27 p.).

5.2 Category 1: DMPs Have a Significant Effect on Participant Behavior

Most experiments ($N = 101/148$ e.; 23/27 p.) tested for differences in participant behavior between a control group which was not exposed to DMPs and treatment group(s) that were exposed to a DMP or multiple DMPs, effectively measuring whether DMPs affected behavior to a statistically significant degree (see Figure 4 for an example). Of those experiments, a majority ($N = 86/101$ e.; 20/27 p.) determined that participant behavior *was changed* when DMPs were experienced, shown by high-level type in Table 4. Given the low subdivided sample counts, no pattern type can be claimed to be more likely to result in statistically significant effect—verifiable by pairwise two-tailed z-tests for two population proportions.

Across all Category 1 experiments, 85% observed that DMPs had a significant effect on participants.[12] Therefore, while it should be

[12]The portion of experiments observing significant effects ($86/101 = 85.1\%$) is different from random chance to a high level of significance ($\alpha = .0005$). The one-tailed one-sample z-test for proportion: $Z = \frac{(p-P)}{\sqrt{\frac{P(1-P)}{n}}} = 7.06 \gg 3.29$, where p is the portion

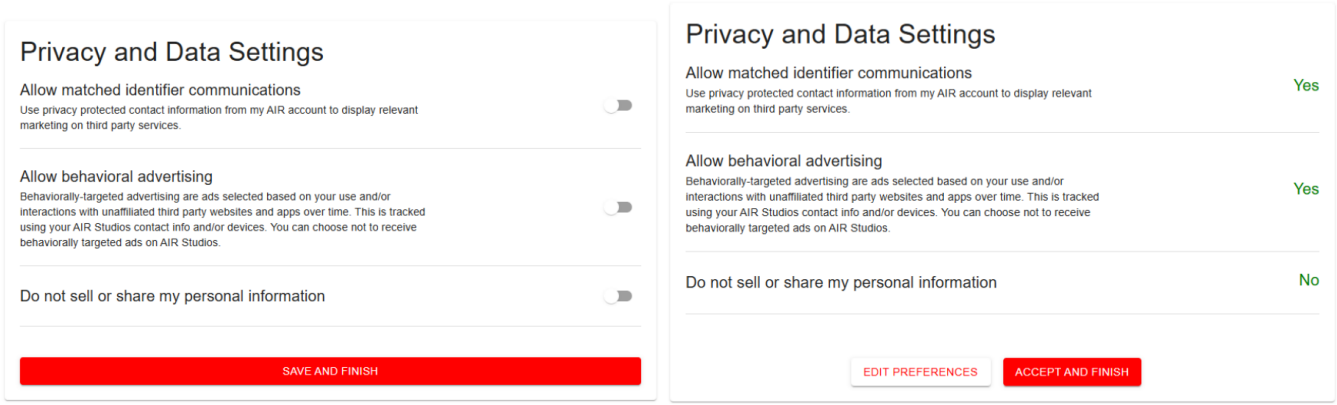

(a) **(b)**

Figure 4: Example from a Category 1 experiment testing how the design of privacy settings affected participant decisions while signing up for a video streaming platform. Kugler et al. observed a statistically significant increase in participants "opting" for data collection when shown an interface with DMPs (b) compared to an interface without (a) [40]. The images are adapted from the original work [40].

Table 4: All 101 experimental units in Category 1. Each circle represents an experimental unit that tested the effects of DMP(s) compared to a control group by high-level DMP type: Interface Interference (II), Social Engineering (SE), Obstruction (Ob), Forced Action (FA), Sneaking (Sn), or a combination of types (Mult.). ●=statistically significant effects. ○=no evidence of significant effects. See A.2 for a mapping of the paper identifiers.

	DMP High-Level Type					
Paper ID	II	SE	Ob	FA	Sn	Mult.
bauer21	●●●●●					
berens24	●●●●●○					
bielova24	●○					
bogliacino24	○					●
boumasims23	●●		●●			●
eucomm22		●●○			●	●●●
gerber23	●					
graßl21	○○		○			
habib22	●●		●●			
koh23		●●●●●				
kugler25	●●			●●●		●●●●●
klütsch23	●					
löschner23					●	
luguri21	●●●○	●●○	●		●●	●●
ma22	○○					
machuletz20	●●●●					
mildner25	●					
naheyan24						○
nouwens20	●●					
o'connor21	●●●○					●●
schaffner25				●●●		
utz19	●					
zac25	●●●●●●●	●●●●	●●○○			●
total:	*41/50*	*13/15*	*7/10*	*6/6*	*4/4*	*15/16*

of positive experimental results observed (0.851) and P is the expected portion of positive experimental results according to random chance (0.500).

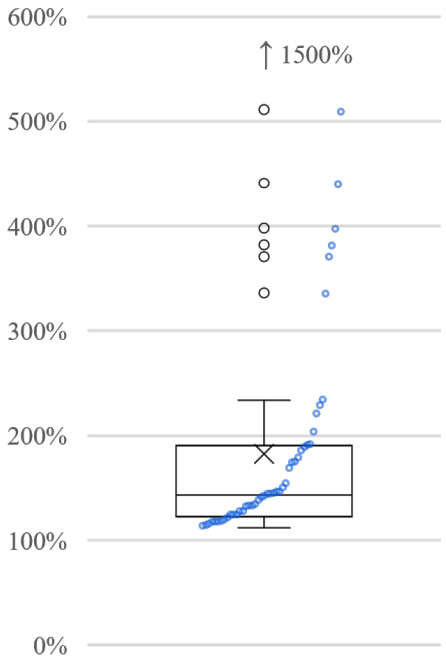

Figure 5: Relative percent increase in the portions of participants experiencing undesirable outcomes in presence of DMPs.

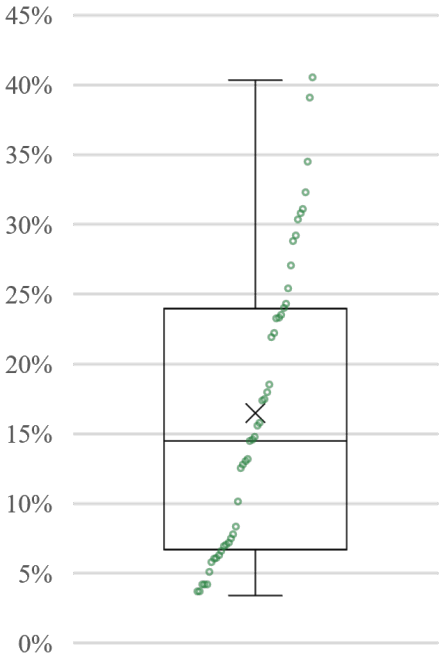

Figure 6: Absolute percent increase in the portions of participants experiencing undesirable outcomes in presence of DMPs.

assumed that some degree of publication bias is present, there remains convincing experimental agreement that DMPs significantly affect participant behavior.

Types of Harms. The Category 1 experiments that observed significant effects found a variety of harms induced from DMPs. Most commonly (N = 52/86 e.; 15/20 p.), participants were found to select privacy un-friendly options as a result of DMPs, such as being more likely to accept unnecessary tracking cookies when pushed by DMPs [e.g., 31]. After privacy-related harms, there is a drop off in the number of studies experimentally measuring other DMP harms. The second most experimentally-measured DMP harm was participants "opting" to purchase goods or services they otherwise would not have (N = 29/86 e.; 3/20 p.)[38, 42, 70]. Other harms included spending more time on a platform than they otherwise would have (N = 3/86 e.; 1/20 p.) [62], downloading apps they otherwise would not (N = 1/86 e.; 1/20 p.) [43], and exhibiting reduced ability to distinguish between real news and advertisements (N = 1/86 e.; 1/20 p.) [41].

Size of Effects. The effects of DMPs were most commonly reported as the difference in the portion of participants experiencing an outcome that they otherwise would not, as concluded by contrast to a control group without the DMPs. For example, Zac et al. showed that the use of Obstruction increased the rate at which participants accepted a product promotion from 48.76% to 61.28% (p<.001) [70]. Of the experiments reporting statistical significance in Category 1, about half reported population means for their control and treatment samples (N = 47/86 e.; 11/20 p.), allowing for the calculation of comparable effect sizes. To control for varied control group means, we calculated both relative and absolute percent increases which are shown in Figure 5 and Figure 6, respectively.

Relative percent increases ranged from 112% to 1500% (mean = 211%, median = 144%). That is, on average, the DMPs tested in these experiments saw participants about 2× more likely to experience an undesirable outcome—e.g., sharing data they otherwise would not have or purchasing an item they otherwise would not have. Absolute percent increases ranged from 3.40% to 40.4% (mean = 16.5%, median = 14.5%). The large variation in effect sizes indicates strong context dependence; that is, the magnitude of DMP effects can vary greatly depending on factors such as the severity of the DMP implemented. However, we found no correlation between the size of the effect and the type(s) of DMPs being studied or the domain of the experiment.

5.3 Category 2: External Interventions That Could (Not) Mitigate DMP Effects

The next most common experimental type (N = 27/148 e.; 5/27 p.) were the experiments that tested whether external interventions could counteract the effect of DMPs. External interventions took the form of: reflective "cool down" interstitials (N = 10/27 e.; 1/5 p.; see Figure 7, e.g.) [43], explicitly priming participants with instructions to select options to maximize privacy (N = 9/27 e.; 1/5 p.) [40], raising the stakes of participant decisions by increasing the cost of a

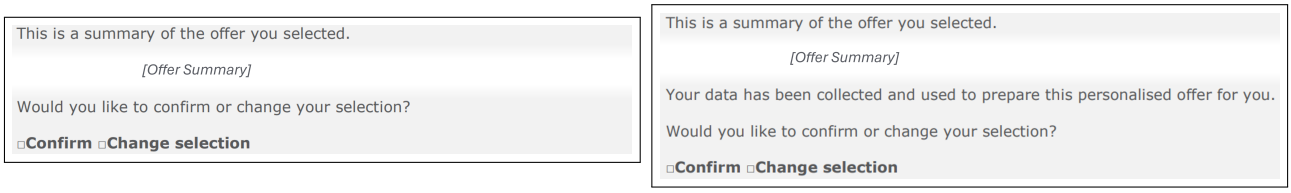

(a) **(b)**

Figure 7: Examples from a Category 2 experiment testing interstitials as interventions. A report by the European Commission tested the effectiveness of two versions of a "cool down" interstitial designed to show participants a summary of the decisions they just had made during a subscription sign-up flow that used DMPs. Both versions allowed participants to change their response. Additionally, one version also disclosed to participants the site's use of a Social Engineering DMP (b) [43]. Neither version of the interstitials was found to reduce the effects of the DMPs. The images are adapted from the original work [43].

Table 5: All experimental units that tested whether external interventions could mitigate the effects of DMPs by high-level DMP type: Interface Interference (II), Social Engineering (SE), Obstruction (Ob), Forced Action (FA), Sneaking (Sn), or a combination of types (Mult.). ●=statistically significant reductions in DMP effects. ○=no evidence of significant reductions. See A.2 for a mapping of the paper identifiers.

	DMP High-Level Type					
Paper ID	II	SE	Ob	FA	Sn	Mult.
eucomm22		○○○○○○○○			○○	
koh23		○○				
kugler25	●●○○			○○		●●○
klütsch23	○					
luguri21						○○○○○
total:	*2/5*	*0/10*	*-*	*0/2*	*0/2*	*2/8*

phony subscription (N = 5/27 e.; 1/5 p.) [42], educating participants on DMPs (N = 2/27 e.; 1/5 p.) [38], and educating participants on the consequences of privacy-related decisions (N = 1/27 e.; 1/5 p.) [37].

As shown in Table 5 only four experiments (N = 4/27 e.; all from [40]) found evidence that the effect of DMPs could be counteracted by external interventions. Specifically, in all four cases, participants were told to maximize their privacy goal with the instructions "choose the most privacy protective options. [...]" when signing up for a service, despite any prior privacy beliefs. Still, the number of experiments (N = 4) where the intervention counteracted the effect of DMPs is fewer than the experiments (from the same work) finding that the exact same intervention did not significantly reduce the DMP's main effects (5/9 e.; also from [40]). Thus, even when participants were explicitly trying to select what they believed to be the most privacy-friendly options, their efforts were largely unsuccessful.

Luguri et al. tested whether increasing the cost of a subscription pushed by DMPs affected subscription rates. We treat this as an intervention in the form of raising the stakes: by making the financial consequences of participants' choices more significant. While not an intervention one would propose as a solution to DMPs, the experiments assess whether participants giving heightened effort were still susceptible to DMPs. As the authors discuss, participants facing higher monetary stakes would (in theory) "be willing to jump over more hurdles in order to save themselves more money" and "be less likely to 'fall' for the DMPs." However, all experiments testing this intervention (N = 5; all from [42]) found that even with cost tripling, DMPs remained just as effective.

There were two versions of the cool down interventions tested (See Figure 7): one that gave participants a summary of their choices (N = 6) and one that disclosed the use of Social Engineering in addition to the summary (N = 4) [43]. Both gave participants the ability to either confirm or change their selections. None of the experiments testing the cool down remedies found significant reductions in the effects of DMPs.

Further, none of the experiments testing participant education (N = 3 e.; 2/5 p.) found statistically significant reductions in the effects of DMPs either. The tested education strategies included: a 7-minute DMP informational video [38], an interactive 'spot-the-dark-pattern' activity [38], and a virtual-assistant-like pop-up that explained the details of the cookie consent choices participants were exposed to [37]. The researchers of these studies discuss possible reasons for the lack of evidence of effective educational interventions being study limitations, such as participant drop-out [38], or cookie decision behavior being already strongly conditioned in Internet users [37].

5.4 Category 3: Whether DMPs Have Additive Effects When Layered is Inconclusive

Some experiments (N = 11/148 e.; 5/27 p.) tested the effects comparing two conditions of DMPs where one group experienced a

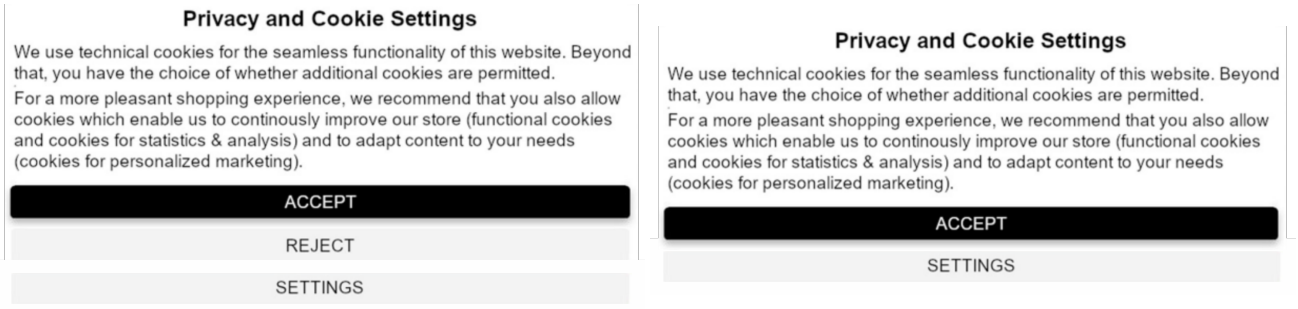

(a) (b)

Figure 8: Example of a Category 3 experiment. Mager et al. tested the difference between two cookie consent interfaces, one with just Interface Interference (a) and one with the same Interface Interference as well as added Obstruction (b) [46]. They found a statistically significant difference in user consent between the two versions. The images are adapted from the original work [46].

(a) (b)

Figure 9: Example of a Category 4 experiment. Berens et al. tested whether embedding a cookie opt-out hyperlink at the end (a) or middle (b) of a consent notice—two distinct instances of Interface Interference—affected participant consent rates [7]. They found no evidence of divergent behavior between the two groups. The images are adapted from the original work [7].

subset of the DMPs of a different group, effectively measuring the marginal effect of additional DMPs (See Figure 8 for an example).

Category 3 experiments form a small subset of the overall literature reviewed and the strength and generalizability of these findings is relatively limited. Three of these studies (N = 4/11 e.; 3/5 p.) [8, 11, 46] found evidence that adding an additional DMP on top of existing DMPs further changed participant behavior. The other studies (N = 7/11 e.; 2/5 p.) [40, 43] found no evidence of increased effect from additional DMP exposure. Thus, while DMPs do effectively change participant behavior, more research is needed to confirm whether they have diminishing marginal effects when multiple are used in conjunction.

5.5 Category 4: Comparisons Between DMP Types are Implementation Dependent

The remaining experiments (N = 9/148 e.; 6/27 p.) tested behavioral differences between multiple experimental conditions where one cannot be considered a subset of another, effectively measuring the comparative effect between different DMPs or sets of DMPs (See Figure 9 for an example). Results of these experiments were mixed and fail to provide overall conclusions on the comparative effects of DMPs. As with Category 3 experiments, these experiments make up a small subset of the overall literature reviewed. As such, findings in this section should be weighed accordingly.

Some Category 4 experiments (N = 6/9 e.; 5/6 p.) showed that certain DMP instances were more effective than others. Results of these experiments showed: (i) that certain instances of Interface Interference and Obstruction more effectively increased participant data sharing than others (N = 3/6 e.; 2/5 p.) [7, 16], (ii) Limited Time Messaging was more effective at guiding participant purchase decisions than three other Social Engineering DMP types (N = 1/6 e.; 1/5 p.) [38], (iii) Negative Framing more strongly resulted in participant's privacy-unfriendly behavior than Positive Framing (N = 1/6 e.; 1/5 p.) [44], and (iv) the change of one instance of Bad Defaults to separate, more severe instances had a significant effect on participant erroneous spending (N = 1/6 e.; 1/5 p.) [57].[13]

[13]While the authors call this a single DMP—"Dark Defaults", referring to the change itself, rather than the comparison of two different instances of DMPs—, such a temporal interpretation does not currently fit into the leading DMPs ontology.

Table 6: All papers that examined the potential mediating role of personal characteristics on the effects of DMP. ●=significant correlation between all measures of the attribute and the strength of the DMP effect. ○=no evidence of attribute interaction with DMPs for all measures. ◐=some but not all measures indicated a correlation. $^{+}$ = a positive correlation with DMP effect strength. $^{-}$ = a negative correlation with DMP effect strength. Some papers found correlations in both directions depending on the measure used ($^{+/-}$).

	Personal Characteristic																
Paper ID	Age	Gender	Education	Income	Country	Domain Knowledge	Privacy Attitudes	Tech. Skills	Risk Aversion	Patience	Trust	Self Esteem	Authoritarianism	Political Orientation	Social Class	ADHD	Financial Literacy
boumasims23		○			○			○									
eucomm22									○	●$^{-}$	●$^{+}$						
fernandez21							●$^{-}$	●$^{-}$									
gerber23						○	◐$^{-}$	○				○					
habib22	○	○															
koh23	●$^{+}$	○								○							
kugler25	○		◐$^{-}$					◐$^{-}$					◐$^{+}$	○	○		
luguri21			◐$^{-}$														
machuletz20							●$^{-}$										
mildner25																◐$^{+/-}$	
naheyan24						◐$^{-}$											
posner23	○	○	○	○													
zac25	◐$^{+/-}$		○	○													◐$^{+}$

However, Category 4 experimental results were not always significant (N = 3/9 e.; 2/6 p.). Two experiments found that participants experiencing varied versions of Forced Action and Interface Interference did not statistically differ from one another [40]. Lastly, one experiment found that slightly different instances of the same DMP were found to have no statistical difference [7]. Ultimately, more experiments need to be conducted in order to make conclusive claims about which DMPs are more likely to cause harm. However, results from these initial studies indicate that specific implementation details (such as DMP complexity, severity, timing, etc.) may be more important than DMP type alone in predicting harm.

5.6 Category 5: The Role of Personal Characteristics

A number of papers (13/27) reported testing for correlations between the effects of DMPs and users' personal characteristics—ranging from demographics (e.g., age and gender) to personal attributes (e.g., technology affinity and political affiliation). These are studies that test whether the presence of a personal characteristic affected the likelihood or strength of a DMP's effect. For example, after Lupiáñez-Villanueva et al. observed how DMPs led consumers to make choices that they would not have made otherwise (recall Table 4), they further tested whether the effects of DMPs were mediated by three self-reported personal characteristics [43]. They found that the effects of DMPs were: (i) not correlated with *risk aversion*, (ii) negatively correlated with *patience*, and (iii) positively correlated with *trust*. That is, DMP effects were constant no matter how participants reported their aversion to risk but were higher for participants that reported lower levels of patience or higher levels of trust. Table 6 summarizes all Category 5 results. As discussed in §4.2, this analysis is conducted at the *paper-level* rather than the *experiment-level* since authors' correlation assertions commonly relied on multiple paper-wide measures.

As shown, many correlations are muddied by inconsistency. Most (9/15) of the papers that identified correlations between DMP effects and personal characteristics also found lack of correlations (or even correlations in the opposite direction) when using other measures. For example, Zac et al. found that older participants were more likely to select a product pushed by Interface Interference (a positive correlation) but less likely to select the same product if it were pushed with Obstruction instead (a negative correlation) [70]. They also found that age did not interact with participants' willingness to actually complete the payment—beyond just selecting the product—when the product was pushed with Obstruction or Social Engineering (no correlations).

In summary, experimental evidence that personal characteristics can mediate the effects of DMPs is sparse (but nonetheless present). This supports claims that differences in personal characteristics are less important than the existence or severity of DMPs. Of course, more research is warranted for conclusivity, especially when considering the specific personal characteristics for which there are still few studies. However, some attribute interactions stand out as approaching early experimental agreement. For example, of all four works analyzing gender [12, 31, 38, 57], none have found significant interactive effects with DMPs. Further, negative correlations between privacy concerns and DMP effects were found in all three papers that tested for it [8, 25, 45]—but not consistently. Additional research is needed to reliably identify which—if any—user groups are consistently more vulnerable to different classes of DMPs.

6 Discussion

6.1 Backing Regulatory Discussions with Empirical Agreement

Given the strong evidence in our review that DMPs result in user harm (§ 5.2) and the lack of strong evidence that external interventions mitigate the effects of DMPs (§ 5.3), we recommend that DMP scholars and policymakers should confront the fact that consumer "self-help" [40] or other "aftermarket tools" may not solve the problem. Even embedded platform tools may not mitigate harms: For example, TikTok's embedded time management tool reduced teens' average daily engagement from about 108.5 minutes by only about 1% [4]. This supports the argument that, rather than solely focusing on an increased public awareness of or tolerance for DMPs, the optimal strategy to safeguard consumers from DMPs likely requires removal of DMPs from the digital landscape, a pursuit well-suited for regulatory bodies and standard-setting organizations.

Regulatory discussions are already underway [20, 23, 56]. Our review adds weight to these discussions by uncovering agreement within DMPs experimental scholarship regarding their significant effects and resilience to interventions, while recognizing that certain DMPs, especially Forced Action and Sneaking, have not been studied as thoroughly.

6.2 Untangling Tensions and Addressing Experimental Gaps

Despite the number of experiments we reviewed and the conclusions reached that DMPs commonly have an effect on consumer behavior, we find that there remain some important open questions.

For one, the literature offers only preliminary results on two key issues: (1) the relative strengths of different DMP types (§ 5.5), and (2) the extent to which combining ("stacking") multiple DMPs amplifies consumer harm (§ 5.4). We recommend future experimental work address these gaps.

There are also dimensions of DMPs that we observed to be overlooked by experimentalists. Most experiments thus far focus on privacy settings, followed by e-commerce. Other possible harms put forth by Mathur et al.'s landmark DMPs work [48] include harms to individual autonomy and societal welfare. Such non-material harms (as Santos et al. call them [61]) related to attention, autonomy, and critical thinking are quite under-represented in the corpus of conducted experiments. While there has been some progress in attempting to isolate and measure societal harms from DMPs [57]—such as their impact on fair markets or cultural perceptions—much work remains to be done, as these effects are still not fully understood and remain difficult to measure in controlled settings.

Other open questions involve the dominant temporal lens of most experimental work. The bulk of studies address short-term, immediate behavioral responses to isolated DMP exposures. However, emerging scholarship suggests that DMPs may have effects that persist or even compound over time and repeated interactions. For instance, sudden changes to an interface can be used to exploit inertia or muscle memory in repeat users [39], and exposure to DMPs in a cookie interfaces can change participant behavior on future, non-dark-pattern cookie interfaces [10]. We did not identify any other work that tested how the effects of DMPs may last beyond the interaction at-hand, potentially indicating an insufficiency in existing experimental approaches for capturing consequences across extended timescales. We encourage researchers to examine long-term and non-material harms of DMPs, as these remain underexplored. Closing these gaps will benefit both future studies and effective policy-making.

6.3 Evidentiary Standards and the Role of Experiments

Understanding the evidentiary basis of DMP research is crucial, especially as regulatory and legal bodies often rely on scientific findings. Recent work has analyzed the range of research methodologies used specifically in the DMPs literature, detailing the advantages and limitations of each [30]. Controlled experiments are typically lauded for their ability to demonstrate causality and offer quantifiable measures of behavioral impact, but they often trade off ecological validity and can be cost prohibitive at scale. In contrast, methods such as observational studies, qualitative content analysis, and surveys can provide broader contextual understanding and surface non-obvious occurrences.

Notably, legal and regulatory decisions around DMPs have historically drawn on a variety of evidence types [30], including non-experimental methods like expert evaluations and telemetry data analysis. Thus, experiments are not the sole source of actionable evidence; policy and enforcement have often relied on a triangulation of methods. In sum, developing a robust, actionable understanding of DMP harms requires integrating the rigor of experimental studies with the breadth of insights from observational and qualitative research.

7 Conclusion

Through a systematic literature review, our study puts forth that there is compelling field-wide evidence that DMPs significantly alter user behavior. Studies included in our analysis predominantly indicate that DMPs lead to undesirable outcomes, such as privacy invasions or avoidable purchases. Further, our review indicates that, despite some efforts to counteract the effects of DMPs through external interventions like privacy goal-setting and user education, DMPs remain effective—pointing to the need for regulatory efforts that remove DMPs altogether rather than relying on consumer education. This is further supported by the notion that DMPs seem

to similarly affect all users. Our review also underscores the need for more experiments to be conducted before concluding whether DMPs have a compounding effect or whether certain DMPs are more pernicious than others. Gaps in DMP research remain, such as expanding experimental results for less-studied harms and understanding DMPs long-term or society-wide impacts. Given that DMPs have significant effects on user behavior, we urge HCI researchers and policymakers to build on this foundation to address both the immediate harms and the persistent, evolving impact of DMPs in the digital landscape.

References

[1] Alessandro Acquisti. 2009. Nudging Privacy: The Behavioral Economics of Personal Information. *IEEE Security & Privacy* 7, 6 (2009), 82–85. doi:10.1109/MSP.2009.163

[2] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM Comput. Surv.* 50, 3, Article 44 (Aug. 2017), 41 pages. doi:10.1145/3054926

[3] Sanju Ahuja and Jyoti Kumar. 2022. Conceptualizations of user autonomy within the normative evaluation of dark patterns. *Ethics and Information Technology* 24, 4 (2022), 52.

[4] Bobby Allyn, Sylvia Goodman, and Dara Kerr. 2024. TikTok executives know about app's effect on teens, lawsuit documents allege. https://www.npr.org/2024/10/11/g-s1-27676/tiktok-redacted-documents-in-teen-safety-lawsuit-revealed

[5] Dan Ariely and Simon Jones. 2008. *Predictably irrational.* HarperCollins New York.

[6] Jan M Bauer, Regitze Bergstrøm, and Rune Foss-Madsen. 2021. Are you sure, you want a cookie?–The effects of choice architecture on users' decisions about sharing private online data. *Computers in Human behavior* 120 (2021), 106729.

[7] Benjamin Maximilian Berens, Mark Bohlender, Heike Dietmann, Chiara Krisam, Oksana Kulyk, and Melanie Volkamer. 2024. Cookie disclaimers: Dark patterns and lack of transparency. *Computers & Security* 136 (2024), 103507.

[8] Carlos Bermejo Fernandez, Dimitris Chatzopoulos, Dimitrios Papadopoulos, and Pan Hui. 2021. This Website Uses Nudging: MTurk Workers' Behaviour on Cookie Consent Notices. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 346 (Oct. 2021), 22 pages. doi:10.1145/3476087

[9] Nataliia Bielova. 2023. A survey of user studies as evidence for dark patterns in consent banners. https://www-sop.inria.fr/members/Nataliia.Bielova/papers/BIEL_CNIL_LINC_2023.pdf

[10] Nataliia Bielova, Laura Litvine, Anysia Nguyen, Mariam Chammat, Vincent Toubiana, and Estelle Hary. 2024. The effect of design patterns on (present and future) cookie consent decisions. In *33rd USENIX Security Symposium (USENIX Security 24)*. 2813–2830.

[11] Francesco Bogliacino, Leonardo Pejsachowicz, Giovanni Liva, and Francisco Lupiáñez-Villanueva. 2024. Testing for manipulation: Experimental evidence on dark patterns. *Available at SSRN 4755295* (2024).

[12] Elijah Robert Bouma-Sims, Megan Li, Yanzi Lin, Adia Sakura-Lemessy, Alexandra Nisenoff, Ellie Young, Eleanor Birrell, Lorrie Faith Cranor, and Hana Habib. 2023. A US-UK usability evaluation of consent management platform cookie consent interface design on desktop and mobile. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–36.

[13] Martin Brenncke. 2024. Regulating dark patterns. *Notre Dame J. Int'l Comp. L.* 14 (2024), 39.

[14] Harry Brignull. 2011. Dark Patterns: Deception vs. Honesty in UI Design. https://alistapart.com/article/dark-patterns-deception-vs.-honesty-in-ui-design/

[15] Corina Cara. 2019. Dark patterns in the media: A systematic review. *Network Intelligence Studies* 7, 14 (2019), 105–113.

[16] Servicio Nacional del Consumidor. 2022. Policy Paper on Cookies Consent Requests: Experimental Evidence of Privacy by Default and Dark Patterns On Consumer Privacy Decision Making. *https://www.sernac.gob.cl/portal/619/w3-article-64969.html* (2022). Translated:https://icpen.org/sites/default/files/2022-05/SERNAC_Policy_Paper_Cookies_Experiment.pdf

[17] Linda Di Geronimo, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. UI Dark Patterns and Where to Find Them: A Study on Mobile Applications and User Perception. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–14. https://dl.acm.org/doi/10.1145/3313831.3376600

[18] The Diversity and Inclusion Council. 2025. Words Matter. https://web.archive.org/web/20250921075524/https://www.acm.org/diversity-inclusion/words-matter Internet Archive | Wayback Machine.

[19] Mateusz Dubiel, Anastasia Sergeeva, and Luis A. Leiva. 2024. Impact of Voice Fidelity on Decision Making: A Potential Dark Pattern?. In *Proceedings of the 29th International Conference on Intelligent User Interfaces* (Greenville, SC, USA) *(IUI '24)*. Association for Computing Machinery, New York, NY, USA, 181–194. doi:10.1145/3640543.3645202

[20] European Data Protection Board 2022. Guidelines 03/2022 on deceptive design patterns in social media platform interfaces: how to recognise and avoid them. https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-032022-deceptive-design-patterns-social-media_en

[21] European Union. 2022. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and Amending Directive 2000/31/EC (Digital Services Act). https://eur-lex.europa.eu/eli/reg/2022/2065/oj/eng Official Journal of the European Union.

[22] Madison Fansher, Shruthi Sai Chivukula, and Colin M Gray. 2018. # darkpatterns: Ux practitioner conversations about ethical design. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–6.

[23] FTC Press Release 2022. FTC Report Shows Rise in Sophisticated Dark Patterns Designed to Trick and Trap Consumers. https://www.ftc.gov/news-events/news/press-releases/2022/09/ftc-report-shows-rise-sophisticated-dark-patterns-designed-trick-trap-consumers

[24] FTC Press Release 2023. FTC Finalizes Order Requiring Fortnite maker Epic Games to Pay $245 Million for Tricking Users into Making Unwanted Charges. https://www.ftc.gov/news-events/news/press-releases/2023/03/ftc-finalizes-order-requiring-fortnite-maker-epic-games-pay-245-million-tricking-users-making

[25] Nina Gerber, Alina Stöver, Justin Peschke, and Verena Zimmermann. 2023. Don't accept all and continue: Exploring nudges for more deliberate interaction with tracking consent notices. *ACM Transactions on Computer-Human Interaction* 31, 1 (2023), 1–36.

[26] Paul Graßl, Hanna Schraffenberger, Frederik Zuiderveen Borgesius, and Moniek Buijzen. 2021. Dark and Bright Patterns in Cookie Consent Requests. In *Journal of Digital Social Research*. https://doi.org/10.33621/jdsr.v3i1.54

[27] Colin M Gray, Nataliia Bielova, Cristiana Santos, and Thomas Mildner. 2024. An Ontology of Dark Patterns: Foundations, Definitions, and a Structure for Transdisciplinary Action. (2024). https://dl.acm.org/doi/10.1145/3613904.3642436

[28] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. 2018. The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–14. https://doi.org/10.1145/3173574.3174108

[29] Colin M. Gray, Lorena Sanchez Chamorro, Ike Obi, and Ja-Nae Duane. 2023. Mapping the Landscape of Dark Patterns Scholarship: A Systematic Literature Review. In *Companion Publication of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) *(DIS '23 Companion)*. Association for Computing Machinery, New York, NY, USA, 188–193. doi:10.1145/3563703.3596635

[30] Johanna T. Gunawan, Colin M. Gray, Cristiana Santos, and Nataliia Bielova. 2025. FORTHCOMING: Leveraging Interdisciplinary Methods for Evidence Collection in Enforcement: Dark Patterns as a Case Study. In *Internet Policy Review Special Issue: The Craft of Interdisciplinary Research and Methods in Public Interest Cybersecurity, Privacy, and Digital Rights Governance (IPR'25 Special Issue)*. Internet Policy Review.

[31] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. 2022. "Okay, whatever": An evaluation of cookie consent interfaces. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–27.

[32] Hilda Hadan, Lydia Choong, Leah Zhang-Kennedy, and Lennart E. Nacke. 2024. Deceived by Immersion: A Systematic Analysis of Deceptive Design in Extended Reality. *ACM Comput. Surv.* 56, 10, Article 250 (May 2024), 25 pages. doi:10.1145/3659945

[33] Johanna Herman. 2024. Dark patterns: EU's regulatory efforts. *Security and Privacy* 7, 6 (2024), e441.

[34] Dennis Hummel and Alexander Maedche. 2019. How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioral and Experimental Economics* 80 (2019), 47–58. doi:10.1016/j.socec.2019.03.005

[35] Athina Ioannou, Iis Tussyadiah, Graham Miller, Shujun Li, and Mario Weick. 2021. Privacy nudges for disclosure of personal information: A systematic literature review and meta-analysis. *PloS one* 16, 8 (2021), e0256822.

[36] Ridha Joober, Norbert Schmitz, Lawrence Annable, and Patricia Boksa. 2012. Publication bias: what are the challenges and can they be overcome? 149–152 pages.

[37] Jennifer Klütsch, Christian Böffel, Sophia von Salm-Hoogstraeten, and Sabine J Schlittmeier. 2023. Defeating Dark Patterns: The impact of supporting information on dark patterns and cookie privacy decisions. *Proceedings TecPsy 2023* (2023), 41.

[38] Woon Chee Koh and Yuan Zhi Seah. 2023. Unintended consumption: The effects of four e-commerce dark patterns. *Cleaner and Responsible Consumption* 11 (2023), 100145. doi:10.1016/j.clrc.2023.100145

[39] Mariliza Kontogeorgou, Christof Van Nimwegen, and Almila Akdag Salah. 2023. Illuminating Muscle Memory's Sinister Side: A Social Media Case Study. In

Proceedings of the European Conference on Cognitive Ergonomics 2023 (Swansea, United Kingdom) *(ECCE '23)*. Association for Computing Machinery, New York, NY, USA, Article 7, 4 pages. doi:10.1145/3605655.3605664

[40] Matthew B Kugler, Lior Strahilevitz, Marshini Chetty, and Chirag Mahapatra. 2025. Can Consumers Protect Themselves Against Privacy Dark Patterns? *University of Chicago Coase-Sandor Institute for Law & Economics Research Paper* 25-01 (2025).

[41] Deborah Maria Löschner and Sebastian Pannasch. 2023. Different ways to deceive: Uncovering the psychological effects of the three dark patterns preselection, confirmshaming and disguised ads. In *International Conference on Human-Computer Interaction*. Springer, 62–69.

[42] Jamie Luguri and Lior Jacob Strahilevitz. 2021. Shining a light on dark patterns. *Journal of Legal Analysis* 13, 1 (2021), 43–109.

[43] Francisco Lupiáñez-Villanueva, Alba Boluda, Francesco Bogliacino, Giovanni Liva, Lucie Lechardoy, and Teresa Rodríguez de las Heras Ballell. 2022. *Behavioural study on unfair commercial practices in the digital environment: dark patterns and manipulative personalisation.* Publications Office of the European Union.

[44] Eryn Ma and Eleanor Birrell. 2022. Prospective consent: The effect of framing on cookie consent decisions. In *CHI Conference on human factors in computing systems extended abstracts*. 1–6.

[45] Dominique Machuletz and Rainer Böhme. 2019. Multiple purposes, multiple problems: A user study of consent dialogs after GDPR. *arXiv preprint arXiv:1908.10048* (2019).

[46] Stefan Mager and Johann Kranz. 2021. On the Effectiveness of Overt and Covert Interventions in Influencing Cookie Consent: Field Experimental Evidence.. In *ICIS*.

[47] Arunesh Mathur, Gunes Acar, Michael J Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–32. https://dl.acm.org/doi/10.1145/3359183

[48] Arunesh Mathur, Mihir Kshirsagar, and Jonathan Mayer. 2021. What Makes a Dark Pattern... Dark?: Design Attributes, Normative Considerations, and Measurement Methods. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–18. https://dl.acm.org/doi/10.1145/3411764.3445610

[49] Thomas Mildner, Daniel Fidel, Evropi Stefanidi, Paweł W Woźniak, Rainer Malaka, and Jasmin Niess. 2025. A Comparative Study of How People With and Without ADHD Recognise and Avoid Dark Patterns on Social Media. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–17.

[50] Alberto Monge Roffarello and Luigi De Russis. 2022. Towards Understanding the Dark Patterns That Steal Our Attention. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–7. https://dl.acm.org/doi/10.1145/3491101.3519829

[51] Alberto Monge Roffarello, Kai Lukoff, and Luigi De Russis. 2023. Defining and Identifying Attention Capture Deceptive Designs in Digital Interfaces. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19. https://dl.acm.org/doi/10.1145/3544548.3580729

[52] Tasneem Naheyan and Kiemute Oyibo. 2024. The effect of dark patterns and user knowledge on user experience and decision-making. In *International Conference on Persuasive Technology*. Springer, 190–206.

[53] Arvind Narayanan, Arunesh Mathur, Marshini Chetty, and Mihir Kshirsagar. 2020. Dark Patterns: Past, Present, and Future: The evolution of tricky user interfaces. *Queue* 18, 2 (2020), 67–92. https://dl.acm.org/doi/10.1145/3400899.3400901

[54] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 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* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376321

[55] Sean O'Connor, Ryan Nurwono, Aden Siebel, and Eleanor Birrell. 2021. (Un) clear and (In) conspicuous: The right to opt-out of sale under CCPA. In *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*. 59–72.

[56] OECD. 2022. Dark commercial patterns. *OECD Digital Economy Papers* (2022).

[57] Nathaniel Posner, Andrey Simonov, Kellen Mrkva, and Eric J Johnson. 2023. Dark defaults: How choice architecture steers political campaign donations. *Proceedings of the National Academy of Sciences* 120, 40 (2023), e2218385120.

[58] Marie Potel-Saville and Mathilde Da Rocha. 2023. From dark patterns to fair patterns? Usable taxonomy to contribute solving the issue with countermeasures. In *Annual Privacy Forum*. Springer, 145–165.

[59] Parinda Rahman and Ifeoma Adaji. 2024. Ethics in Persuasive Technologies: A Systematic Literature Review. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia (MUM '24)*. Association for Computing Machinery, New York, NY, USA, 106–118. doi:10.1145/3701571.3701572

[60] Maria Rosala. 2023. Nielsen Norman Group: Deceptive Patterns in UX: How to Recognize and Avoid Them. https://www.nngroup.com/articles/deceptive-patterns/

[61] Cristiana Santos, Viktorija Morozovaite, and Silvia De Conca. 2025. No harm no foul: how harms caused by dark patterns are conceptualised and tackled under EU data protection, consumer and competition laws. *Information & Communications Technology Law* (2025), 1–47.

[62] Brennan Schaffner, Yaretzi Ulloa, Riya Sahni, Jiatong Li, Ava Kim Cohen, Natasha Messier, Lan Gao, and Marshini Chetty. 2025. An experimental study of Netflix use and the effects of autoplay on watching behaviors. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–22.

[63] State of California. 2020. California Privacy Rights Act of 2020 (Proposition 24). https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5 California Civil Code, Title 1.81.5, Sec. 1798.100–1798.199.100.

[64] Daniel Susser and Vincent Grimaldi. 2021. Measuring Automated Influence: Between Empirical Evidence and Ethical Values. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (Virtual Event, USA) *(AIES '21)*. Association for Computing Machinery, New York, NY, USA, 242–253. doi:10.1145/3461702.3462532

[65] Richard H Thaler and Cass R Sunstein. 2021. *Nudge: The final edition.* Penguin.

[66] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 973–990. doi:10.1145/3319535.3354212

[67] Fiona Westin and Sonia Chiasson. 2020. Opt out of privacy or "go home": understanding reluctant privacy behaviours through the FoMO-centric design paradigm. In *Proceedings of the New Security Paradigms Workshop* (San Carlos, Costa Rica) *(NSPW '19)*. Association for Computing Machinery, New York, NY, USA, 57–67. doi:10.1145/3368860.3368865

[68] Lauren E Willis. 2020. Deception by design. *Harv. JL & Tech.* 34 (2020), 115.

[69] Weiwei Yi and Zihao Li. 2025. Mapping the scholarship of the regulation of dark patterns: A systematic review of concepts, regulatory paradigms, and solutions from law and HCI perspectives. *Computer Law & Security Review* 59 (2025), 106225.

[70] Amit Zac, Yu-Chun Huang, Amédée von Moltke, Christopher Decker, and Ariel Ezrachi. 2023. Dark patterns and consumer vulnerability. *Behavioural Public Policy* (2023), 1–50.

A Supplementary Materials

A.1 Search Queries

Google Scholar and ACM DL query: ("dark pattern" OR "dark patterns" OR "dark design" OR "dark designs" OR "deceptive design" OR "deceptive designs" OR "manipulative design" OR "manipulative designs" OR "coercive design" OR "coercive designs" OR "dark default" OR "dark defaults") AND (experiment*). SSRN query: "dark patterns OR dark defaults". ArXiv query "dark patterns".

A.2 Supplementary Tables

Table 7: List of included papers and the number of experimental units per paper. The government reports are bolded.

Paper ID	Exp. Units	Citation
bauer21 [6]	6	Bauer, J. M., Bergstrøm, R., Foss-Madsen, R. (2021). Are you sure, you want a cookie?–The effects of choice architecture on users' decisions about sharing private online data. *Computers in Human behavior*, *120*, 106729.
berens24 [7]	8	Berens, B. M., Bohlender, M., Dietmann, H., Krisam, C., Kulyk, O., Volkamer, M. (2024). Cookie disclaimers: Dark patterns and lack of transparency. *Computers Security*, *136*, 103507.
bielova24 [10]	2	Bielova, N., Litvine, L., Nguyen, A., Chammat, M., Toubiana, V., Hary, E. (2024). The effect of design patterns on (present and future) cookie consent decisions. In *33rd USENIX Security Symposium (USENIX Security 24)* (pp. 2813-2830).
bogliacino24 [11]	3	Bogliacino, F., Pejsachowicz, L., Liva, G., & Lupiáñez-Villanueva, F. (2023). Testing for manipulation: Experimental evidence on dark patterns. *Available at SSRN 4755295*.
boumasims23 [12]	5	Bouma-Sims, E. R., Li, M., Lin, Y., Sakura-Lemessy, A., Nisenoff, A., Young, E., Birrell, E., Cranor, L.F. Habib, H. (2023, April). A US-UK usability evaluation of consent management platform cookie consent interface design on desktop and mobile. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-36).
eucomm22 [43]	20	Lupiáñez-Villanueva, F., Boluda, A., Bogliacino, F., Liva, G., Lechardoy, L., & de las Heras Ballell, T. R. (2022). Behavioural study on unfair commercial practices in the digital environment: dark patterns and manipulative personalisation.*Publications Office of the European Union.*
fernandez21 [8]	2	Bermejo Fernandez, C., Chatzopoulos, D., Papadopoulos, D., Hui, P. (2021). This website uses nudging: Mturk workers' behaviour on cookie consent notices. *Proceedings of the ACM on human-computer interaction*, *5*(CSCW2), 1-22.
gerber23 [25]	1	Gerber, N., Stöver, A., Peschke, J., Zimmermann, V. (2023). Don't accept all and continue: Exploring nudges for more deliberate interaction with tracking consent notices. *ACM Transactions on Computer-Human Interaction*, *31*(1), 1-36.
graßl21 [26]	3	Graßl, P., Schraffenberger, H., Borgesius, F. Z., Buijzen, M. (2021). Dark and Bright Patterns in Cookie Consent Requests. *Journal of Digital Social Research*, *3*(1), 1-38.
habib22 [31]	4	Habib, H., Li, M., Young, E., Cranor, L. (2022, April). "Okay, whatever": An evaluation of cookie consent interfaces. In *Proceedings of the 2022 CHI conference on human factors in computing systems* (pp. 1-27).
koh23 [38]	8	Koh, W. C., Seah, Y. Z. (2023). Unintended consumption: The effects of four e-commerce dark patterns. *Cleaner and Responsible Consumption*, *11*, 100145.
kugler25 [40]	26	Kugler, M. B., Strahilevitz, L., Chetty, M., Mahapatra, C. (2025). Can Consumers Protect Themselves Against Privacy Dark Patterns?. *University of Chicago Coase-Sandor Institute for Law Economics Research Paper*, (25-01).
klütsch23 [37]	2	Klütsch, J., Böffel, C., von Salm-Hoogstraeten, S., Schlittmeier, S. J. (2023). Defeating Dark Patterns: The impact of supporting information on dark patterns and cookie privacy decisions. *Proceedings TecPsy 2023*, 41.

Continued on next page

Table 7: List of included papers and the number of experimental units per paper. The government reports are bolded. (Continued)

löschner23 [41]	1	Löschner, D. M., Pannasch, S. (2023, July). Different ways to deceive: Uncovering the psychological effects of the three dark patterns preselection, confirmshaming and disguised ads. In *International Conference on Human-Computer Interaction* (pp. 62-69). Cham: Springer Nature Switzerland.
luguri21 [42]	17	Luguri, J., Strahilevitz, L. J. (2021). Shining a light on dark patterns. *Journal of Legal Analysis, 13*(1), 43-109.
ma22 [44]	3	Ma, E., Birrell, E. (2022, April). Prospective consent: The effect of framing on cookie consent decisions. In *CHI Conference on human factors in computing systems extended abstracts* (pp. 1-6).
machuletz20 [45]	4	Machuletz, D., Böhme, R. (2020). Multiple Purposes, Multiple Problems: A User Study of Consent Dialogs after GDPR. *Proceedings on Privacy Enhancing Technologies*.
mager21 [46]	1	Mager, S., Kranz, J. (2021). On the Effectiveness of Overt and Covert Interventions in Influencing Cookie Consent: Field Experimental Evidence. In *ICIS*.
mildner25 [49]	1	Mildner, T., Fidel, D., Stefanidi, E., Woźniak, P. W., Malaka, R., & Niess, J. (2025, April). A Comparative Study of How People With and Without ADHD Recognise and Avoid Dark Patterns on Social Media. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1-17).
naheyan24 [52]	1	Naheyan, T., Oyibo, K. (2024, April). The effect of dark patterns and user knowledge on user experience and decision-making. In *International Conference on Persuasive Technology* (pp. 190-206).
nouwens20 [54]	2	Nouwens, M., Liccardi, I., Veale, M., Karger, D., Kagal, L. (2020, April). 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).
o'connor21 [55]	6	O'Connor, S., Nurwono, R., Siebel, A., Birrell, E. (2021, November). (Un) clear and (In) conspicuous: The right to opt-out of sale under CCPA. In *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society* (pp. 59-72).
posner23 [57]	1	Posner, N., Simonov, A., Mrkva, K., & Johnson, E. J. (2023). Dark defaults: How choice architecture steers political campaign donations. *Proceedings of the National Academy of Sciences*, *120*(40), e2218385120.
schaffner25 [62]	3	Schaffner, B., Ulloa, Y., Sahni, R., Li, J., Cohen, A. K., Messier, N., Gao, L., Chetty, M. (2025). An Experimental Study Of Netflix Use and the Effects of Autoplay on Watching Behaviors. *Proceedings of the ACM on Human-Computer Interaction, 9*(2), 1-22.
sernac22 [16]	1	Pavón Mediano, A. Rabanales F., Luengo-Miranda, C. Vergara, G. (2022). Policy Paper on Cookies Consent Requests: Experimental Evidence of Privacy by Default and Dark Patterns On Consumer Privacy Decision Making. From *Servicio Nacionaldel Consumidor.*
utz19 [66]	1	Utz, C., Degeling, M., Fahl, S., Schaub, F., Holz, T. (2019, November). (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).
zac25 [70]	16	Zac, A., Huang, Y. C., von Moltke, A., Decker, C., Ezrachi, A. (2023). Dark patterns and consumer vulnerability. *Behavioural Public Policy*, 1-50.