Matthew B. Kugler,[*] Lior Jacob Strahilevitz,[**] Marshini Chetty,[***] Chirag Mahapatra[****] Yaretzi Ulloa[*****]

Can Consumers Protect Themselves Against Privacy Dark Patterns?

23 U.N.H. L. Rev. 243 (2025)

Abstract. Dark patterns have emerged in the last few years as a major target of legislators and regulators. Dark patterns are online interfaces that manipulate, confuse, or trick consumers into purchasing goods or services that they do not want, or into surrendering personal information that they would prefer to keep private. As new laws and regulations to restrict dark patterns have emerged, skeptics have countered that motivated consumers can and will protect themselves against these manipulative interfaces, making government intervention unnecessary. This debate occurs alongside active legislative and regulatory discussion about whether to prohibit dark patterns in newly enacted comprehensive consumer privacy laws. Our interdisciplinary paper provides experimental evidence showing that consumer self-help is unlikely to fix the dark patterns problem. Several common dark patterns (obstruction, interface interference,

[*] Professor of Law, Northwestern Pritzker School of Law

[**] Sidley Austin Professor of Law, University of Chicago

[***] Associate Professor of Computer Science, University of Chicago

[****] Harvard University, Kennedy School of Government.

[*****] Yale University. The authors thank Courtney Cox, Adam Davidson, Christoph Engel, James Hicks, Andrew Kent, Hajin Kim, Eric Martinez, Jonathan Masur, Martha Nussbaum, Paul Schwartz, Sonja Starr, lecture and workshop attendees at the American Law and Economics Association Annual Meeting at the University of Colorado / Princeton University State Attorney General Forum on Privacy Regulation, participants at the 2025 Privacy Law Scholars Conference, the United Kingdom Competition and Markets Authority, Fordham University Law School, the University of California-Berkeley Law School, and the University of Chicago Law School, the Carl S. Lloyd Faculty Fund at the University of Chicago for research support, and Caroline Cole and Elijah Griesz for research assistance.

preselection, and confusion), which we integrated into the privacy settings for a video-streaming website, remain strikingly effective at manipulating consumers into surrendering private information even when consumers were charged with maximizing their privacy protections and understood that objective. We also provide the first published evidence of the independent potency of "nagging" dark patterns, which pester consumers into agreeing to an undesirable term. These findings strengthen the case for legislation and regulation to address dark patterns. Our paper also highlights the broad popularity of a feature of the recent California Consumer Privacy Act (CCPA), which gives consumers the ability to opt-out of the sale or sharing of their personal information with third parties. As long as consumers see the Do Not Sell option, a super-majority of them will exercise their rights, and a substantial minority will even overcome dark patterns in order to do so.

INTRODUCTION

Regulators around the world have dark patterns in their crosshairs, with a flurry of new regulations passed in the last several years.[1] Dark patterns are user interfaces whose designers knowingly confuse users, make it difficult for users to express their preferences or satisfy their objectives, or manipulate users into taking actions that are inconsistent with their preferences or well-being.[2] Dark patterns typically exploit cognitive biases with the goal of prompting users to purchase goods or services they do not want or surrendering personal information they prefer to keep private.[3]

Examples of dark patterns will be familiar to anyone who navigates the Internet or uses a smartphone, even if the terminology is not. A smartphone app might ask whether a user will authorize push notifications, with the only two response options being "Yes" and "Maybe Later." Then if the user selects "Maybe Later" she will be asked the same question again days later, and again days after that. But once a user clicks on "Yes" the user will never be asked to reconsider this choice; this is a "nagging" dark pattern.[4] Alternatively, someone wishing to make a one-time purchase may find

---

1 *See* Part I.

2 Jamie Luguri & Lior Jacob Strahilevitz, *Shining a Light on Dark Patterns*, 13 J. LEGAL ANAL. 43 (2021); Colin M. Gray, Cristiana Santos, Nataliia Bielova & Thomas Mildner, *An Ontology of Dark Patterns Knowledge: Foundations, Definitions, and a Pathway for Shared Knowledge-Building*, in' CHI 24: PROCEEDINGS OF THE 2024 CHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS (May 2024), available at https://dl.acm.org/doi/10.1145/3613904.3642436; *see also* Cal. Priv. Prot. Agency, CAL. CODE REGS. TIT. 11, § 7004(c) (2023) ("A user interface is a dark pattern if the interface has the effect of substantially subverting or impairing user autonomy, decisionmaking, or choice.").

3 Dark patterns promote impulsive decision-making and exploit cognitive biases such as loss aversion and the sunk-cost fallacy. See Agnieszka Kitkowska, *The Hows and Whys of Dark Patterns: Categorizations and Privacy*, in HUMAN FACTORS IN PRIVACY RESEARCH 173, 189–91 (Nina Gerber, Alina Stover & Karola Marky eds. 2023); Arjun Sharma, *Uncovering Dark Patterns of Persuasive Design* (UI/UX), 11 J. ENG'G DESIGN & ANAL. 1, 4 (2024); Ray Sin et al., *Dark Patterns in Online Shopping: Do They Work and Can Nudges Help Mitigate Impulse Buying?*, BEHAVIOURAL PUB. POL'Y 1, 2 (2022).

4 We rely here on classic taxonomies of dark patterns, such as Colin M. Gray et al., *The Dark (Patterns) Side of UX Design*, PROC. 2018 CHI CON. ON HUM. FACTORS IN COMPUTING SYS., https://dl.acm.org/doi/abs/10.1145/3173574.3174108 [https://perma.cc/9BAL-JPRK]; Arunesh Mathur et al., *Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites*, 3 PROC. ACM ON HUMAN-COMPUTER INTERACTION 1-32 (2019), https://dl.acm.org/doi/10.1145/3359183 [https://perma.cc/GY9L-F5FL]; and Arunesh Mathur, Mihir Kshirsagar & Jonathan Mayer, *What Makes a Dark Pattern . . . Dark? Design Attributes, Normative Considerations, and Measurement Methods*,

themselves subscribing to recurring monthly purchases because that option is selected by default; this is a "preselection" dark pattern. Unappealing aspects of a service may be buried in fine print or a light gray font – an "interface interference" dark pattern. Or opting out of privacy-invasive company practices may needlessly require numerous and time-consuming mouse clicks, an "obstruction" dark pattern. All told, researchers have identified dozens of dark patterns, including the aforementioned ones and other manipulative strategies like making it easy to sign up for a service but hard to cancel, using double negatives to confuse users into making decisions inconsistent with their preferences, or forcing people to agree to obnoxious "confirmshaming" statements like "I like wasting money" when they wish to opt out of receiving a small discount in exchange for having their purchases tracked.[5]

Whereas a few years ago the academic literature on dark patterns was quite sparse, scores of new papers on dark patterns are now appearing every year, largely in computer science but increasingly in legal scholarship as well. In that time, a scholarly consensus has emerged that many variants of dark patterns are highly effective at convincing consumers to purchase goods or services they do not want, and that they are proliferating online despite the efforts of regulators and legislators to keep them in check.[6] Indeed, restricting dark patterns has become something of a cat-and-mouse game. As the problem has grown and the manipulative potential has become increasingly evident, a debate has emerged in the literature over whether regulation is appropriate. Skeptics of regulatory and enforcement efforts argue that the ubiquity of dark patterns makes them manageable. In the skeptics' view, consumers are becoming increasingly familiar with dark patterns and so, over time, they may become increasingly adept at overcoming them.[7] Perhaps in response to this sentiment, American jurisdictions are splitting over the question of whether comprehensive privacy laws, enacted at the state level, need to include

---

PROC. 2021 CHI CONF. ON HUM. FACTORS IN COMPUTING SYS., https://dl.acm.org/doi/abs/10.1145/3411764.3445610 [https://perma.cc/X5GV-NFVF].

5 For a catalogue of dark patterns with colorful examples, *see* Harry Brignull et al., *Types of Deceptive Pattern*, DECEPTIVE PATTERNS, https://www.deceptive.design/types [https://perma.cc/76F9-LUTK] (last visited Feb. 27, 2025).

6 *See* Part II.

7 *See infra* notes 58–66 and accompanying text.

specific provisions prohibiting the use of dark patterns to thwart consumers' efforts to exercise their opt-out rights. At present, roughly two-thirds of these jurisdictions include dark pattern prohibitions in their state privacy laws and one third do not.[8]

Our paper squarely addresses the related arguments that consumers can defend themselves against dark patterns and that legal prohibitions on dark patterns are therefore unnecessary. It is the first paper to show that substantial numbers of consumers are unable to resist the manipulative effects of dark patterns even when they are instructed to do so and understand their objective. We asked a census-weighted sample of American adults to select the most privacy protective settings on a mock video-streaming website we built from scratch, and then we placed dark pattern obstacles in their paths when they navigated through their privacy settings. While dark patterns had a larger effect on research subjects instructed to choose the settings they would normally select than on those subjects directed to choose the most privacy-protective options, the effects of dark patterns in the user interfaces on both groups were still significant and substantial. This evidence suggests that even when consumers are trying to protect their privacy, the kinds of dark patterns they regularly encounter online can confuse, manipulate, or pester them into surrendering private information and privacy rights. In short, our experiment gave consumers a clear goal of adopting privacy protective settings on a website, consumers fully understood that goal and tried to protect their privacy, and then dark patterns caused many of them to fail at their task. Furthermore, our paper is also the first to show the independent efficacy of nagging dark patterns, which repeatedly pester consumers to consent if they initially decline or threaten to keep asking for permission if consumers do not relent the first time. These results strengthen the case for muscular regulatory and statutory interventions.

I. LAWS RELATING TO DARK PATTERNS

Concern about dark patterns has grown in the past few years, with legislators enacting new statutes, regulators promulgating novel regulations, and enforcement agencies using laws new and old to sue companies that are employing dark patterns to manipulate consumers. On August 12, 2024, the White House announced a major

[8] *See infra* notes 35–37 and accompanying text.

new "Time is Money" initiative, with several key federal agencies coalescing around an agenda that places dark patterns in the federal government's crosshairs and various other efforts to thwart consumer preferences and waste consumers' time.[9] This initiative includes the Federal Trade Commission "click to cancel" regulation, which was finalized on October 16, 2024, and requires companies to make cancelling a subscription as easy and quick as signing up for one.[10] The Time is Money initiative also includes other regulations targeting needless friction that obstructs consumers who wish to cancel or modify their services and subscriptions.[11]

Domestically, California was the first mover, enacting restrictions on the use of dark patterns in the California Privacy Rights Act (CPRA).[12] Under that law, which amended the California Consumer Privacy Act (CCPA), California consumers have the right to opt out of the sale or sharing of their personal information. The statute explicitly prohibits apps and web sites subject to the law from using dark patterns to discourage such opt-outs, defining dark patterns as user interfaces "designed or manipulated with the substantial effect of impairing user autonomy, decision-making, or choice."[13] The law further provided that consumer consent obtained via dark patterns would be ineffective as a matter of law,[14] and it authorized the California Privacy Protection Agency to develop more detailed regulations defining what behaviors amount to dark patterns.[15]

---

[9] *See Fact Sheet: Biden Harris Administration Launches New Effort to Crack Down on Everyday Headaches and Hassles that Waste Americans' Time and Money*, THE WHITE HOUSE (Aug. 12, 2024), https://web.archive.org/web/20250116072209/https://www.whitehouse.gov/briefing-room/statements-releases/2024/08/12/fact-sheet-biden-harris-administration-launches-new-effort-to-crack-down-on-everyday-headaches-and-hassles-that-waste-americans-time-and-money/ [https://perma.cc/SDS6-YXAS]. This document was later removed by the Trump Administration.

[10] *See* FTC Negative Option Rule, 16 C.F.R. pt. 425 (2024).

[11] *See Fact Sheet: Biden Harris Administration Launches New Effort to Crack Down on Everyday Headaches and Hassles that Waste Americans' Time and Money*, *supra* note 9.

[12] *See* CAL. CIV. CODE § 1798.140(l) (West 2025); Jeremy Merkel, *Dark Patterns Come to Light in California Data Privacy Laws*, NAT. L. REV. (July 2, 2021), https://natlawreview.com/article/dark-patterns-come-to-light-california-data-privacy-laws [https://perma.cc/887V-N26K] ("As is often the case with consumer protection, California is the first state to regulate dark patterns").

[13] CAL. CIV. CODE § 1798.140(l) (West 2025).

[14] *Id.* § 1798.140(h).

[15] *Id.* § 1798.140(l).

Those California regulations were promulgated in 2023, becoming the first detailed substantive laws to target dark patterns in e-commerce.[16] Under § 7004 of these regulations, the design of user interfaces that consumers wishing to opt-out would encounter need to adhere to several principles. More precisely, they must be (a) easy to understand, (b) symmetrical in choice, (c) non-confusing, (d) non-impairing and non-interfering, and (e) easy to execute.[17] While some of these design principles are straightforward, like the "easy to understand" and "non-confusing" rules that guard against intentional consumer confusion, others required extensive elaboration in the regulations.[18] For example, symmetry in choice requires that companies make it no harder for consumers to exercise the privacy protective choice than the less-protective option.[19] Thus, requiring a consumer to click through seven screens to protect their privacy but only two to waive their privacy rights would be a dark pattern. The symmetry in choice framework also prohibits websites or app developers from giving users a loaded choice between "Yes" and "Ask Me Later" when a request to sell or share personal information is presented, and requiring consumers to choose between "Accept all" and "Preferences" (requiring more clicks to give a more nuanced answer) is an impermissible dark pattern too.[20] The former exemplifies a nagging dark pattern in that a user who does not wish to share personal information can anticipate that she will be asked repeatedly for consent, whereas a user who clicks "Yes" can expect the prompts to stop, with the consumer having given the response that the user interface designers were hoping for. The latter is a kind of obstruction dark pattern, where the user interface designer is interposing needless obstacles in the path of a consumer who does not wish to accept all tracking cookies. Similarly, a user interface designer who created an "Accept All" tracking cookies button would have to also offer a "Decline All" option.[21]

---

16 *Cf.* Merkel, *supra* note 12.

17 CAL. CODE REGS. 11, § 7004 (2025).

18 *Id.* The "easy to understand" provision in the regulation simply says, "The methods shall use language that is easy for consumers to read and understand." Many of the other provisions contain multiple examples of acceptable and unacceptable practices.

19 *Id.*

20 *Id.* § 7004(a)(2)(C).

21 *Id.*

Impairment/interface interference dark pattern regulations prohibit efforts to secure consent that is not "freely given, specific, informed, and unambiguous."[22] For example, the regulations prohibit forcing consumers to click through disruptive screens to effectuate their opt-out rights, and they also bar requiring consumers to consent to incompatible uses of their data that are bundled with desirable features.[23] Finally, the easy execution design principle prohibits interface designers from requiring consumers to scroll through a large wall of text to locate the links and fields required to complete the opt-out process. It also requires them to repair and maintain opt-out links so that they are neither circular nor broken, and it prohibits user interfaces that force users seeking to opt out to wait unnecessarily while an opt-out request is processed.[24] The easy execution rules, in short, target varieties of obstruction dark patterns. Notably, the California regulations do not target all recognized types of dark patterns. For example, various forms of social-engineering dark patterns are not addressed by the regulation, such as confirmshaming (the use of emotionally manipulative language that forces consumers to affirm language they disagree with in order to protect their privacy) or Social Proof (creating a bandwagon effect that taps into consumers' propensity to conform to the apparent behavioral norm).[25] In recent months, California has begun enforcing its prohibitions on dark patterns, with the California Privacy Protection Agency advising the technology industry that it means business.[26] New empirical work done by a team of researchers, including two authors of this paper, finds that the California restrictions on dark patterns have been moderately effective, though some new dark patterns have emerged to exploit loopholes in the regulatory regime.[27] Because California looms so large in the American economy, and in the technology sector in particular, the CCPA has a very significant extraterritorial impact. Recent research

---

22 *Id.* § 7004(a)(4).

23 *Id.* § 7004(a)(4)(B).

24 *Id.* § 7004(a)(5).

25 *See* Gray et al., *supra* note 2, at 1.

26 *See* Cal. Priv. Prot. Agency Enf't Div., *Applying Data Minimization to Consumer Requests* (Apr. 2, 2024), https://cppa.ca.gov/pdf/enfadvisory202401.pdf [https://perma.cc/6NKH-FXVD].

27 Van Tran, Aarushi Mehrotra, Rayna Sharma, Marshini Chetty, Nick Feamster, Jens Frankenreiter & Lior Strahilevitz, *Dark Patterns in the Opt-Out Process and Compliance with the California Consumer Privacy Act*, Proc. 2025 CHI Conf. on Hum. Factors in Computing Sys. (2025), https://dl.acm.org/doi/full/10.1145/3706598.3714138 [https://perma.cc/2KNJ-Q76A].

suggests that in some respects the CCPA has become a *de facto* national privacy law, influencing the design of websites and apps consumers encounter nationwide.[28]

Other states, such as Colorado and Connecticut, have enacted similarly explicit prohibitions on dark patterns as a means of securing consent to process personal information.[29] Colorado regulations require that consent interface options be presented to consumers "in a symmetrical way that does not impose unequal weight or focus on one available choice over another such that a [c]onsumer's ability to consent is impaired or subverted."[30] The regulations provide examples of dark patterns such as making an "I accept" button that is larger or in a more prominent style than the "I do not accept" button, or offering an "accept all" button without offering a "reject all" button.[31] The regulations also prohibit the use of emotionally manipulative confirmshaming, treating inaction as consent, default terms that are less protective of privacy, obstruction, and deceptive or intentionally confusing language.[32]

With efforts to enact comprehensive federal privacy legislation having stalled,[33] numerous states have entered the void and enacted their own comprehensive privacy laws.[34] Such states face an important fork in the road. Should their own

---

[28] Van Hong Tran, Aarushi Mehhrotra, Marshini Chetty, Nick Feamster, Jens Frankenreiter & Lior Strahilevitz, *Measuring Compliance with the California Consumer Privacy Act over Space and Time*, PRO. 2024 CHI CONF. ON HUM. FACTORS IN COMPUTING SYS. (May 11, 2024), https://dl.acm.org/doi/full/10.1145/3613904.3642597 [https://perma.cc/3BHA-ZKV4].

[29] Colorado's Privacy Act and Connecticut's Data Privacy Act says that consumer agreement obtained via dark patterns does not constitute consent. COLO. REV. STAT. § 6-1-1303(5) (2025); CONN. GEN. STAT. § 42-515(7) (2025). Existing private law causes of action at the state level may supplement these statutory efforts. *See* Gregory M. Dickinson, *Privately Policing Dark Patterns*, 57 GA. L. REV. 1633 (2023).

[30] COLO. CODE REGS § 904-3:7.09(A)(1) (2025).

[31] *Id.*

[32] *Id.* § 904-3:7.09(A)(7).

[33] Comprehensive federal privacy legislation has been repeatedly introduced over the past several years. It most recently failed in the summer of 2024. *See, e.g.*, Catherine Stupp, *Patchwork of State Privacy Laws Remains After Latest Failed Bid for Federal Law*, WALL ST. J. (Aug. 27, 2024), https://www.wsj.com/articles/patchwork-of-state-privacy-laws-remains-after-latest-failed-bid-for-federal-law-2a1a020d [https://perma.cc/A3KE-UCGH].

[34] *See, e.g.*, Tony Foley, *Five New Comprehensive State Privacy Laws Take Effect: What Businesses Need to Know*, LAW.COM (Feb. 11, 2025), https://www.law.com/legaltechnews/2025/02/11/five-new-

comprehensive privacy laws follow California, Colorado, and Connecticut in prohibiting dark patterns? Or should they leave dark patterns unmentioned in the state privacy statutes? By our count, among the nineteen states that have enacted such laws to date, twelve (California, Colorado, Connecticut, Delaware, Maryland, Minnesota, Montana, Nebraska, New Hampshire, New Jersey, Rhode Island, and Texas) have adopted legal provisions specifying that consumer consent is ineffective if obtained via the use of dark patterns,[35] one state (Oregon) does not mention "dark patterns" explicitly but says that consent cannot be achieved through mechanisms that impair consumer autonomy,[36] and six states (Indiana, Iowa, Kentucky, Tennessee, Utah, and Virginia) neither mention nor reference dark patterns in their consumer privacy laws.[37] While these six states all allow consumers to opt out of the sale of their personal information and to access their personal information,[38] websites and apps may be able to use dark patterns to make the effective exercise of those privacy rights quite cumbersome.

Despite congressional inaction and recent state legislative efforts, there are important developments at the federal level too. The U.S. Federal Trade Commission (FTC), in addition to proposing the aforementioned "click to cancel" rule, has also

---

comprehensive-state-privacy-laws-take-effect-what-businesses-need-to-know/ [https://perma.cc/5AFJ-HLKJ].

[35] Calif. Consumer Privacy Act, CAL. CIV. CODE § 1798.185(a)(20)(c)(iii) (West 2025); Colorado Privacy Act, COLO REV. STAT § 6-1-1303(5) (2025); Connecticut Personal Data Privacy and Online Marketing Act, Conn. Act No. 22-15 §1(6) (Reg. Sess.); Delaware Personal Data Privacy Act, DEL CODE ANN. tit. 6, § 12D-102(7) (2025); Maryland Online Data Privacy Act, MD. CODE ANN., COM. LAW § 14-4601(G)(3) (effective Oct. 1, 2025); Minnesota Consumer Data Privacy Act, MINN. STAT. § 325O.02(f) (2025); Montana Consumer Data Privacy Act, S.B. 384 §2(5)(b)(iii). 68th Leg., Reg. Sess. (Mont. 2023); Nebraska Data Privacy Act LB 1074, NEB. REV. STAT. § 2(6)(b)(iii) (2025); S.B. 255 § 1, Reg. Sess. (N.H. 2023); N.J. REV. STAT. § 56:8-166.4 (2025); R.I. Data Transparency and Privacy Protection Act, R.I. GEN. LAWS § 6-48.1-2(6) (effective Jan. 1, 2026); Texas Data Privacy and Security Act, TEX. BUS. & COM. CODE § 541.001(6)(C) (West 2025).

[36] Oregon Consumer Privacy Act, SB 619 § 1(6), 82nd Gen. Leg. Assemb., Reg. Sess. (Or. 2023).

[37] *See* Consumer Data Protection, S.B. 5, 123rd Gen. Assemb., Reg. Sess. (Ind. 2023); An Act Relating to Consumer Data Protection, Providing Civil Penalties, and Including Effective Date Provisions, S.F. 262, 90th Gen. Assemb. (Iowa 2023); An Act Relating to Consumer Data Privacy and Making an Appropriation Therfor, H.B. 15, Reg. Sess. (Ky. 2024); Tennessee Information Protection Act, S.B. 73, H.B. 1181 113th Gen. Assemb. (Tenn. 2023); Utah Consumer Privacy Act, S.B. 227, 2022 Gen. Sess. (Ut. 2022); Consumer Data Protection Act, VA. CODE ANN. § 59.1-575 (2021).

[38] *See* Oregon Consumer Privacy Act, SB 619 § 1(6).

begun using its existing authority under Section 5 of the FTC Act to sue firms that utilize dark patterns, arguing that dark patterns are unfair or deceptive practices in trade.[39] Pursuing those arguments resulted in substantial settlements, including a $520 million settlement with Epic Games, the makers of Fortnite, over that game's use of dark patterns,[40] a $100 million settlement with Vonage over that company's use of dark patterns,[41] and an $18.5 million settlement with Publishers Clearing House, a sweepstakes marketing entity.[42] The F.T.C. has secured important preliminary litigation victories as well, especially in its suit over dark patterns in Amazon Prime.[43] The Commission has supplemented these enforcement actions with a staff report providing guidance on what dark patterns it deems especially problematic.[44]

The European Union's Digital Services Act,[45] which went into effect fully in February 2024, contains a broader prohibition on dark patterns than the CCPA, and the prohibition is applicable to all conduct by online platforms,[46] not just invocations

---

[39] Luguri & Strahilevitz, *supra* note 2, at 83; Lindsay Wilson, Note, *Is There a Light at the End of the Dark-Pattern Tunnel?*. 91 GEO. WASH. L. REV. 1048, 1052 (2023).

[40] Epic Games, Inc., F.T.C. No. 192-3203, 2023 WL 2609446 (Mar. 13, 2023). Of this amount, Epic paid $275 million in penalties for privacy violations, and another $245 million in refunds. Kimberly A. Berger et al., *More than Child's Play: $520 Million FTC Settlement Signals Risks for Digital Platforms*, NAT. L. REV. (Jan. 27, 2025), available at https://natlawreview.com/article/more-childs-play-520-million-ftc-settlement-signals-risks-digital-platforms#google_vignette [https://perma.cc/5FUA-X79D].

[41] *FTC Action against Vonage Results in $100 Million to Customers Trapped by Illegal Dark Patterns and Junk Fees when Trying to Cancel Service*, F.T.C. (Nov. 3, 2022), https://www.ftc.gov/news-events/news/press-releases/2022/11/ftc-action-against-vonage-results-100-million-customers-trapped-illegal-dark-patterns-junk-fees-when-trying-cancel-service [https://perma.cc/PG6Q-VF6C]. The FTC also invoked the federal Restore Online Shoppers' Confidence Act (ROSCA), 15 U.S.C. § 8403, in this enforcement action. Complaint, F.T.C. v. Vonage Holdings Corp., No. 3:22-cv-6435, 2022 WL 16833021, at *1 (D.N.J. Nov. 3 2022).

[42] *See* Publishers Clearing House, F.T.C. No. 182-3145, 2023 WL 4349342 (June 26, 2023).

[43] *See* F.T.C. v. Amazon.com, Inc., 735 F. Supp.3d 1297 (W.D. Wash. 2024).

[44] *Bringing Dark Patterns to Light*, F.T.C. BUREAU OF CONSUMER PROT. (Sep. 2022), https://www.ftc.gov/system/files/ftc_gov/pdf/P214800%20Dark%20Patterns%20Report%209.14.2022%20-%20FINAL.pdf [https://perma.cc/Z6QD-RV25].

[45] 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) [hereinafter Digital Services Act].

[46] *Id.* at ¶ 67.

of CCPA rights like the right to delete, the right to know, and the right to opt out of the sale or sharing of personal information.[47] The EU law defines dark patterns as "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 or decisions."[48] The law includes examples of dark patterns, such as "presenting choices in a non-neutral manner" and "giving more prominence to certain choices through visual, auditory, or other components, when asking the recipient of the service for a decision."[49] In addition to prohibiting these interface interference approaches, the law's text explicitly prohibits nagging dark patterns, asymmetry in choice/roach motels (whereby it's easy to sign up for a service but hard to cancel), obstruction dark patterns, and pre-selected defaults that "are very difficult to change."[50]

While the Digital Services Act contains Europe's broadest prohibitions on dark patterns, other laws restrict the practices as well. For example, the European Data Protection Board interprets Europe's General Data Protection Regulation (GDPR) to prohibit dark patterns on social media platform interfaces.[51] Though the terminology of the Data Protection Board is somewhat idiosyncratic, these guidelines similarly target obstruction, nagging, interface interference, and confusion dark patterns.[52] The Data Protection Board grounds its dark patterns prohibitions in several GDPR provisions, including Article 5's data minimization and transparency requirements, Article 4 and 7's consent provisions, Article 12's requirement that communication to data privacy subjects be intelligible and easily accessible, and Article 25's data protection by design and default provisions.[53] Europe's Unfair Commercial Practice Directive provides another set of restrictive

---

47 Cal. Code Regs. Tit. 11, § 7004 (2025).

48 Digital Services Act, *supra* note 45.

49 *Id.*

50 *Id.*

51 Eur. Data Prot. Bd., *Guidelines 03/2022 on Deceptive Design Patterns in Social Medial Platform Interfaces: How to Recognise and Avoid Them Version 2.0* (Feb. 24, 2023), https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-032022-deceptive-design-patterns-social-media_en [https://perma.cc/LZJ3-8VTH].

52 *Id.*

53 *Id.* at 4, 11–12.

rules that regulators are turning to in the fight against dark patterns.[54] The United Kingdom recently enacted the Digital Markets, Consumers, and Competition Act of 2024, a law that prohibits various kinds of dark patterns, such as efforts to obstruct subscription cancellations, efforts to confuse consumers who are trying to cancel subscriptions, reminder-free automatic renewals, and drip pricing.[55] Legislators and regulators in Canada are also considering following the lead of American and European lawmakers.[56]

To summarize, then, in just the last few years there has been a flurry of new laws and regulations targeting dark patterns, as well as muscular enforcement actions by agencies using broad language in older vintage laws.[57] Of course, following President Trump's inauguration, his appointment of Andrew N. Ferguson as Chairman, and the President's efforts to fire the F.T.C.'s Democratic Commissioners,[58] enforcement priorities in Washington, D.C. maybe be in flux.

II. KEY POLICY DEBATES

One of the key debates over dark patterns concerns the extent to which consumers are able to defend themselves against firms that use dark patterns. Scholars who argue for legislation and regulation targeting dark patterns generally suggest that the market itself will not deter firms from employing dark patterns. The most widely-cited argument along these lines comes from Luguri and Strahilevitz, who provided empirical evidence suggesting that when firms employ dark patterns in subscription sales they can prompt a substantial increase in the percentage of consumers who purchase subscriptions without generating a significant backlash

---

[54] *See* Mark Leiser, *Illuminating Manipulative Design: From "Dark Patterns" to Information Asymmetry and the Repression of Free Choice under the Unfair Commercial Practices Directive*, 34 Loy. Consumer L. Rev. 484 (2023).

[55] Digital Markets, Competition and Consumers Act 2024, c. 13 §§ 230, 258–261 (UK).

[56] See, e.g., Matthew Gaulton et al., *Approaches to Regulating Privacy Dark Patterns* (Western Univ. Fac. of Info. & Media Stud. Working Paper No. 383, 2024), https://ir.lib.uwo.ca/fimspub/383/ [https://perma.cc/286W-DFEP].

[57] *See, e.g.*, the FTC's use of ROSCA in the Vonage case. Complaint, F.T.C.v. Vonage Holdings Corp., No. 3:22-cv-6435, 2022 WL 16833021, at *1 (D.N.J. Nov. 3, 2022).

[58] *See* David McCabe, *The Two Democrats Trump Fired from the F.T.C. Sue over Their Dismissals*, N.Y. Times, Mar. 27, 2025.

from consumers.[59] Luguri and Strahilevitz did find that extremely blatant and persistent dark patterns generated such a customer backlash, but companies can get away with utilizing a few dark patterns without making consumers less willing to use their services in the future.[60] These experimental results are buttressed by observational evidence about the proliferation of dark patterns online.[61] If dark patterns do generate a substantial backlash from consumers, it is hard to understand why they would be proliferating in e-commerce rather than dying off.

On the other side of the debate are scholars such as Gus Hurwitz, as well as industry advocates, who question whether dark patterns regulations are needed.[62] Hurwitz argues that dark pattern regulation is only appropriate to address the most egregious cases.[63] Considering dark pattern regulations, he posits:

> A better approach to addressing concerns like this is to rely on competition. Customers are generally keenly aware of design issues. There is little better way to drive customers away from a product than for it to have an awkward, cumbersome, or 'unfriendly' interface. When firms are able to compete, and

---

[59] Luguri & Strahilevitz, *supra* note 2, at 67–70, 79–81.

[60] *Id.* at 67–68. Another study found that experimental subjects in an online shopping simulation who were opposed to a dark pattern where premium shipping was added surreptitiously to their shopping carts at checkout had a more negative attitude towards the shopping site. Janis Witte et al., *Consequences of User Manipulation Through Dark Patterns*, PROC. 2023 ICIS CONF. ON HUM. TECH. INTERACTION, https://aisel.aisnet.org/icis2023/hti/hti/5/ [https://perma.cc/9K7Q-QDQ6]. The same study found that a scarcity dark pattern, where customers were told that supply was quite limited, did not generate such negative sentiment. *Id.* While the paper's measure of willingness to do business with a web site again is less direct than the one employed in Luguri & Strahilevitz, the results suggest that there is not a market penalty for firms that employ relatively mild dark patterns, though there may be a disincentive to employ especially obnoxious dark patterns.

[61] *See, e.g.*, Mathur et al., *supra* note 4.

[62] Justin (Gus) Hurwitz, *Designing a Pattern, Darkly*, 22 N.C. J. L. & TECH. 57 (2020). See also Katri Nousianinen & Catalina Perdomo Ortega, *Dark Patterns in Law and Economics Framework*, 36 LOY. CONS. L. REV. 90, 108-115 (2023) (discussing whether market failures exist to justify regulatory interventions with respect to dark patterns). Industry concerns voiced during the notice and comment process for the CPRA's dark patterns regulations similarly articulated the view that firms needed to be able to try to convince consumers of the value of waiving their privacy rights, and worried that the regulations would prevent this sort of persuasion. *See, e.g.*, Letter from Digit. Advert. All. to Lisa B. Kim at 00114-00115 (Oct.2020), https://www.oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-written-comm-3rd-15-day-period.pdf [https://perma.cc/2G89-YBRY].

[63] Hurwitz, *supra* note 62, at 95–100.

> especially where there is evidence that firms compete, regulation over design elements or design decisions is likely undesirable except in the rarest of cases of overtly intentional or exceptionally harmful design patterns . . . . Given the complexity of design, there is reason to prefer to rely on the marketplace to address the concerns raised by dark patterns – particularly given that this market-based approach appears to be working.[64]

These scholars suggest that self-regulation by industry is the best way to deal with most dark patterns[65] or that consumers can learn to defeat dark patterns over time,[66] raising the possibility that the dark patterns problem can be remedied through firms' restraint or consumer self-help.[67] Some scholars are developing automated ways to detect dark patterns, which might lend themselves to software-based solutions to a software-based problem.[68] Moreover, Hurwitz speculates that

---

[64] *Id.* at 89–90, 93.

[65] *See, e.g.*, *id.* at 101.

[66] Tasneem Naheyan & Kiemute Oyino, *The Effect of Dark Patterns and User Knowledge on User Experience and Decision-Making*, PROC. 2024 19TH INT'L CONF., PERSUASIVE TECH. at 190, 203 (finding, in an MTurk study of 211 Canadians, that users of a hypothetical streaming service who were knowledgeable about preselection and confirmshaming dark patterns were better able to resist them); Dominique Kelly & Jacquelyn Burkell, *Identifying and Responding to Privacy Dark Patterns* (Western Fac. of Info. & Media Stud., Working Paper, 2024), https://ir.lib.uwo.ca/fimspub/385/ [https://perma.cc/WDA2-NC3Z].

[67] Studies reach divergent results on whether less educated subjects are more vulnerable to dark patterns, a factor that might shed light on the effectiveness of self-help defenses. *Compare* Amit Zac et al., *Dark Patterns and Online Consumer Vulnerability*, BEHAVIOURAL PUBLIC POL'Y (forthcoming) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4547964 [https://perma.cc/RBX6-A4DE] (finding only weak evidence that less educated British consumers are more vulnerable to dark patterns), *and* Francesco Bogliacino et al., *Testing for Manipulation: Experimental Evidence on Dark Patterns* (Working Paper, 2024), https://osf.io/preprints/socarxiv/sqt3j/ [https://perma.cc/3AZ5-SH8Q] (finding that average consumers may be more vulnerable to dark patterns than less educated consumers), *with* Luguri & Strahilevitz, *supra* note 2 (finding that less educated consumers are more vulnerable to dark patterns), *and* Rebecca Abbott et al., *The Role of Dark Pattern Stimuli and Personality in Online Impulse Shopping: An Application of S-O-R Theory*, 22 J. CONSUMER BEHAV. 1311 (2023) (same).

[68] *See, e.g.*, Jieshan Chen et al., *Unveiling the Tricks: Automated Detection of Dark Pattern*, PROC. 2023 ACM SYMP. ON USER INTERFACE SOFTWARE & TECH. (Oct. 2023), https://dl.acm.org/doi/10.1145/3586183.3606783 [https://perma.cc/V62K-YQNE]; Jordan Donnelly et al., *"Be a Pattern for the World": The Development of a Dark Patterns Detection Tool to Prevent Online User Loss*, PROC. 2022 ETHICOMP CONF. ON ETHICAL & SOC. ISSUES IN COMMC'N TECH., Sept. 9, 2019, at 577 https://arrow.tudublin.ie/ascnetart/3/ [https://perma.cc/2W8L-TFFM]; Ryan Wood,

dark patterns will be more effective in situations where the stakes are relatively low for consumers,[69] perhaps indicating that consumers' vulnerability to dark patterns is a consequence of their lack of motivation. On that logic, maybe consumers can defeat dark patterns when they are focused on doing so and actually trying. As Hurwitz sees it, this phenomenon would weaken the case for legislative and regulatory intervention.[70] The half-dozen states that have so far decided not to regulate dark patterns while enacting comprehensive consumer privacy laws could point to these scholarly defenses of a laissez faire approach in justifying their legislative design choices.[71]

III. GAPS IN THE EXISTING EXPERIMENTAL LITERATURE

There is growing experimental literature on the effectiveness of dark patterns, though the literature is surprisingly narrow in its scope. The first experimental papers on dark patterns were a large-scale experiment by Christine Utz and co-authors in 2019,[72] which examined the effects of changing the position, content, and details of cookie consent notices on visitors to a German-language e-commerce website,[73] a small-scale experiment on a convenience sample by Midas Nouwens and co-authors in 2020,[74] which examined dark patterns on a cookie Consent Management Platform (CMP), and a large-scale experiment by Luguri and Strahilevitz on a census-weighted sample, which was first posted to SSRN in 2019

Understanding the Impact of Dark Pattern Detection on On-line Users (July 17, 2023) (M.S. Thesis, Va. Tech. Univ.) http://hdl.handle.net/10919/115787 [https://perma.cc/JQ82-8GKY].

69 Hurwitz, *supra* note 62 at 91.

70 *Id.* ("If the effect is only limited to low-value transactions, the impact on consumers may not be sufficient to justify regulation that may or may not prove effective. Accordingly, if the concern is that firms use dark patterns to extract small, additional revenue from a large number of consumers that may be particularly at-risk of exploitation, caution [about enforcing laws meant to combat dark patterns] may be particularly warranted.")

71 *See supra* text accompanying note 37.

72 Christine Utz et al., *(Un)informed Consent : Studying GDPR Consent Notices in the Field*, Proc. 2019 SIGSAC Conf. on Comput. & Commc'n Sec., https://arxiv.org/pdf/1909.02638 [https://perma.cc/FSX5-BCZV].

73 *Id.* at 5.

74 Midas Nouwens et al., *Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence*, Proc. 2020 CHI Conf. on Hum. Factors in Computing Sys., https://dl.acm.org/doi/10.1145/3313831.3376321 [https://perma.cc/F22F-SX3S].

and published in final form in 2021.[75] The Luguri and Strahilevitz paper studied the effectiveness of dark patterns strategies to try to manipulate users into signing up for a subscription to an identity-theft protection plan that most of them did not want.[76] All three papers showed that several dark pattern strategies were effective in manipulating consumers into making choices that were inconsistent with their preferences, though each also found that some interface design choices did not have significant effects on choices (the content of disclosures in Utz et al., the notification style of the CMP in Nouwens et al., and the use of urgency/scarcity messages in Luguri and Strahilevitz).[77]

In the years that followed, several other large-scale experiments on the effectiveness of dark patterns have confirmed the core finding of these first three papers, which is that dark patterns are highly effective at prompting consumers to make choices that are inconsistent with their preferences or interests. Yet all these large-scale studies of the effects of dark patterns focus on decision environments similar to that studied by Luguri and Strahilevitz in 2019: the decision to purchase or otherwise obtain a good or service. For example, Zac et al. (2023) find, in a well-done and large-sample experiment, strong evidence that dark patterns are quite effective at manipulating British consumers to purchase an investment product.[78] Zac et al.'s experiment involved having consumers evaluate a fictitious investment website, then pitching them on investing through pop-up notifications that either used dark patterns or did not.[79] Bogliacino et al. (2023) find in a very large-scale study of consumers in six countries that dark patterns are highly effective at manipulating consumer decisions about whether to subscribe to an entertainment website, even when no deception is involved in the experimental design.[80] Furth-Matzkin and Kricheli-Katz (2022) found in an unpublished, large-scale online experiment that dark patterns are quite effective in nudging American consumers towards products they do not prefer, and that dark patterns are especially effective

---

75 Luguri & Strahilevitz, *supra* note 2, at 43.

76 *Id.* at 46.

77 Nouwens et al., *supra* note 74, at 1; Utz et al., *supra* note 72, at 10; Luguri & Strahilevitz, *supra* note 2, at 75.

78 Zac et al., *supra* note 67.

79 *Id.*

80 Bogliacino et al., *supra* note 67, at 1.

at manipulating consumers who are pressed for time, have lower income, and are members of racial minority groups.[81] Furth-Matzkin and Kricheli-Katz's study was based on a gift card lottery in which experimental subjects participated.[82] Sin et al. (2022) find dark patterns to be effective in increasing consumers' propensity to make impulsive purchases in an online shopping experiment.[83]

The evidence of dark pattern effectiveness in these purchase and purchase-like settings is not limited to online and laboratory experiments. New observational evidence from a natural experiment study of dark patterns on political campaign contribution websites confirms that preselection dark patterns are both highly effective and harm consumers thanks to unintentional donations. Posner et al. (2023) studied what happened when some political campaigns began including a pre-checked box on donation pages that would cause recurring monthly campaign contributions (of, say, $20) to become automatically recurring *weekly* $20 donations by default.[84] Donors who wished to avoid contributing weekly would need to uncheck the box.[85] According to Federal Election Commission data, immediately after some campaigns began including the pre-checked box they saw a threefold increase in recurring weekly donations, but donors did not decrease the amounts of each donation as their giving was shifted from once a month to once a week.[86] Donors did start requesting refunds at much higher rates after campaigns made this change, indicating that many of the consumers unknowingly or mistakenly increased their donations above what they were comfortable giving because the weekly recurrence box was checked by default, and learned they had been fooled upon reviewing their credit card statements.[87] Political campaigns that refused to implement the dark pattern did not see increases in weekly contributions, indicating

---

81 MEIRAV FURTH-MATZKIN & TAMAR KRICHELI-KATZ, THE DARK SIDE OF DARK PATTERNS (2022 draft) (on file with author).

82 *Id.*

83 Sin et al., *supra* note 3, at 1.

84 Nathaniel Posner et al., *Dark Defaults: How Choice Architecture Steers Political Campaign Donations*, 120 PROC. OF NAT. ACADEMY OF SCIS. 1–6 (2023), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10556642/pdf/pnas.202218385.pdf [https://perma.cc/3QRK-EJPK].

85 *Id.* at 1.

86 *Id.* at 4.

87 *Id.* at 3.

that the dark pattern itself, rather than extrinsic circumstances, explained the change in donor behavior.[88]

This focus on purchase decisions raises the question of whether dark patterns are effective at prompting consumers to surrender private information in the same way that they can manipulate consumers into surrendering cash. To date, almost all the experimental studies that pose this question follow Utz et al. and Nouwens et al. in studying the choice architecture of CMP and cookie consent interfaces.[89] These studies generally replicate the results of the earliest work on cookie consent interfaces, but drill down to identify particularly effective or ineffective dark patterns.[90] Privacy dark pattern studies not arising from CMP contexts are few and far between, and generally use non-representative samples or convenience samples.[91] For example, Anaraky et al. (2023) find, in a small-*n* MTurk study, that dark patterns are quite effective at prompting consumers to disclose their private information via photo tagging on Facebook, and they find that older consumers are more vulnerable to dark patterns than younger ones.[92]

To the best of our knowledge, no published experimental research study finds that dark patterns are ineffective at manipulating consumers. This is the case even though a well-designed, contrarian study finding null effects from dark patterns would garner significant attention from scholars, not to mention a rousing welcome from industry lobbyists and large law firms' defense counsel, who would like to be

---

88 *Id.* at 2.

89 A helpful overview of this literature is Nataliia Bielova et al., *Two Worlds Apart! Closing the Gap Between Regulating EU Consent and User Studies*, 37 HARV. J. L. & TECH. 1295, 1306–11 (2023). Some studies try to explore design alternatives that are friendlier to consumers. See, e.g., Hana Habib et al., *"Okay, whatever": An Evaluation of Cookie Consent Interfaces*, 2022 CHI CONF. ON HUM. FACTORS IN COMPUT. SYS.

90 See, e.g., Nataliia Bielova et al., *The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions*, 2024 USENIX SEC. SYMP., https://www.usenix.org/system/files/sec23winter-prepub-365-bielova.pdf [https://perma.cc/NKE8-J3JG] (analyzing the effects of highlighted text, text content that underscores the salience of tracking, and the persistence of dark pattern effects over time).

91 *See, e.g.*, Naheyan & Oyino, *supra* note 66, at 192–93.

92 Reza Anaraky et al., *Older and Younger Adults are Influenced Differently by Dark Pattern Designs* (Working Paper 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4581552 [https://perma.cc/9Z69-84GZ].

able to downplay the significance of the threat posed by dark patterns in ongoing legislative debates and litigation.

At the same time, the dark patterns literature examining their effect on privacy settings is quite narrow. Nearly all the work on privacy and dark patterns follows the path blazed by work first done in 2019 and 2020, which focused on consent management platforms and cookie choices, an issue that looms especially large in Europe because of GDPR. This focus on cookies and consent leaves important questions about the efficacy of dark patterns on other privacy settings unanswered. Given that dark pattern regulation, as shown in Part I, is largely a privacy topic at this stage, this gap is disconcerting.

Expanding the experimental privacy literature on dark pattern effectiveness beyond the CMP context is one important contribution in this study. Consumers typically encounter cookie consent interfaces at a moment when they are impatient. They are trying to navigate to a particular website, likely to get particular information needed for some other task, and the cookie consent interface is the frustrating obstacle standing in their way before they can access the content or services they want.[93] Also, the choice of whether one website, among thousands, sets a cookie may seem irrelevant to a busy consumer. As a result, they often want to get past the cookie consent screen as quickly as they can, so it is no wonder that dark patterns have proven quite effective in that circumstance. Yet many consumer choices happen in other contexts, for example when a consumer is first starting a new subscription. Consumers setting up a new account and beginning a new service likely expect to make a series of consequential choices — what subscription to choose, what content to preference, and what data to share. Our experiment confronts consumers with this different (but still familiar and real-world) decision-making environment. Participants in our study are signing up for a hypothetical subscription and trying to make decisions about managing their private information. Our research subjects also know that these settings are not an obstacle standing between them and the show they wish to binge-watch. They have been told,

---

[93] *See* Hai Le & Sirisha Sharon Nethala, Beyond the Banner: Understanding the Impact of Cookie Consent Interfaces on User Data Privacy Choices (May 2024) (Master's thesis, Lund University) https://lup.lub.lu.se/luur/download?func=downloadFile&recordOId=9159294&fileOId=9159297 [https://perma.cc/L5Q7-DWK9]; *Cf.* Utz et al., *supra* note 72, at 982.

and understand, that they are helping us beta-test a new Netflix-like streaming platform, not a live platform that they will use to watch shows and films that evening. We are also compensating them for their time. These dynamics mitigate the kind of frustration that makes many consumers vulnerable to CMP-cookie dark patterns. In that sense, our experiment prompts consumers to be much more attentive to, and considerate of, their privacy choices and settings than a typical CMP interaction does.

There also remains the important policy question: can the effectiveness of dark patterns be negated if consumers are actively seeking to resist them? Some scholars wonder whether consumers can defeat many dark patterns if they are motivated to do so and focused on the task.[94] Other scholars argue that the reams of new dark patterns regulations are, at best, unnecessary.[95] Despite the immense practical import of this hypothesis, there was, prior to our study, no direct test of it.

IV. TESTING THE EFFECTIVENESS OF GOALS IN MITIGATING THE POWER OF DARK PATTERNS

This study is aimed at answering two questions: do dark patterns influence choices consumers make when selecting privacy settings and, if so, does that influence persist even when people are actively seeking to protect their privacy? Participants were asked to go through a sign-up process for a fake video streaming website. The sign-up process gave participants a series of six privacy choices that were highly similar to the kinds of privacy choices consumers are generally asked to make as they sign up for new accounts on commercial websites. Some of the participants were exposed to dark patterns during the sign-up process, and some were not. We received approval from the Institutional Review Boards (IRBs) at Northwestern University and the University of Chicago before running these experiments.

---

[94] *See* Luguri & Strahilevitz, *supra* note 2; Zac et al., *supra* note 67; Bogliacino et al., *supra* note 67; Utz et al., *supra* note 72; Nouwens et al., *supra* note 74; Furth-Matzkin & Kricheli-Katz, *supra* note 81; Sin et al., *supra* note 3; Posner et al., *supra* note 84.

[95] *See supra* text accompanying note 62.

Before beginning the sign-up process, half the participants were assigned a privacy goal. Rather than being asked to sign up for the website as they normally would, they were instead told to make privacy protective choices throughout the process. The other half were told to pick the privacy settings they normally would when signing up for a new video streaming service.

The design was therefore a 2 x 3. Participants either had a privacy goal or did not (2), and participants were either exposed to a control condition without dark patterns; a condition with several different kinds of dark patterns; or a condition with some dark patterns, but specifically without a nagging dark pattern (3). Consistent with best practices in empirical scholarship, we pre-registered the experiment and our hypotheses with aspredicted.org.[96] The main hypotheses we test in this paper were all pre-registered.

A. Procedure

More than 1,700 American adult participants were recruited from the CloudResearch service Connect, which is a professionally managed panel. Entering the study, participants found themselves in a Qualtrics survey. The survey stated that researchers were conducting a usability study for a new video streaming website which we called AIR Studios. Specifically, we claimed to be interested in beta-testing our sign-up process.[97] Participants completed a variety of standard demographic questions as well as questions about their current streaming subscriptions and current use of streaming sites. We also administered a right-wing authoritarianism scale and a technology skills scale, both justified as part of an effort to understand attitudes, experiences, and content preferences. These scales are described below.

Before sending participants to the fake video streaming website, we gave them instructions. All participants were instructed to proceed through the sign-up process as if they were enrolling as actual users. They were also all assured, "You will

[96] *See* Chetty et al., *Dark Patterns, User Goals, and Privacy Settings - An Experimental Study (#175828)* (May 20, 2024, 5:02 PM), https://aspredicted.org/6m7j-g2d3.pdf [https://perma.cc/2QMF-L8A5].

[97] "We are designing a new video streaming site and are interested in how different kinds of people experience the signup process. We will begin with a series of demographic and personality style questions before redirecting you to the website to test it."

NOT be asked for payment information and will NOT actually be signing up. This is just a test of the website."

The instructions continued with the goal manipulation. Half the participants, selected at random, were instructed, "As you go through the process, we would like you to choose the **most privacy protective options**. Whatever your prior beliefs about privacy, see if you can choose the options that would best protect your privacy as a user of the site." The other half were instructed, "As you go through the process, we would like you to choose **whatever options you think you normally would**. Go through the signup process as if you were really signing up for a website." Participants then received an attention-check question on a subsequent page asking them what their goal was during the sign-up process. The correct answer for the privacy protective condition was "To choose the option that most protects my privacy." The correct answer for the normal condition was "To choose the option I normally would." We excluded respondents who answered this question incorrectly.

The website itself was hosted separately from Qualtrics and its flow was modeled after the sign-up process then used by Netflix. The graphics and interface of the website were designed to look similar to those of real video streaming websites. For example, actual movie titles and art appeared on one of the later screens. We made a conscious choice not to make the experiment incentive-compatible, for example by rewarding subjects who overcame dark patterns with a larger payment. We did that for external validity reasons. Namely, when consumers make decisions about privacy settings in the real world, there typically are not immediate financial payoffs or penalties that stem from their choices. In prior research on the use of dark patterns to generate subscription revenue, such as Luguri & Strahilevitz,[98] it was important that research subjects believed they had financial skin in the game. Here, by contrast, the experiment is rendered more informative about how actual consumers behave by virtue of the complex, long-run, and often financially ambiguous consequences of our subjects' choices about privacy settings.

Upon entering the website, participants saw a screen with a realistic video streaming site backdrop that asked for their email address. See Figure 1. They were

[98] Luguri & Strahilevitz, *supra* note 2, at 43.

then prompted to create a password, though neither the password nor the email address was saved. This was done to mirror the sign-up experiences that participants may have had as customers on other websites.

Figure 1: Landing Page

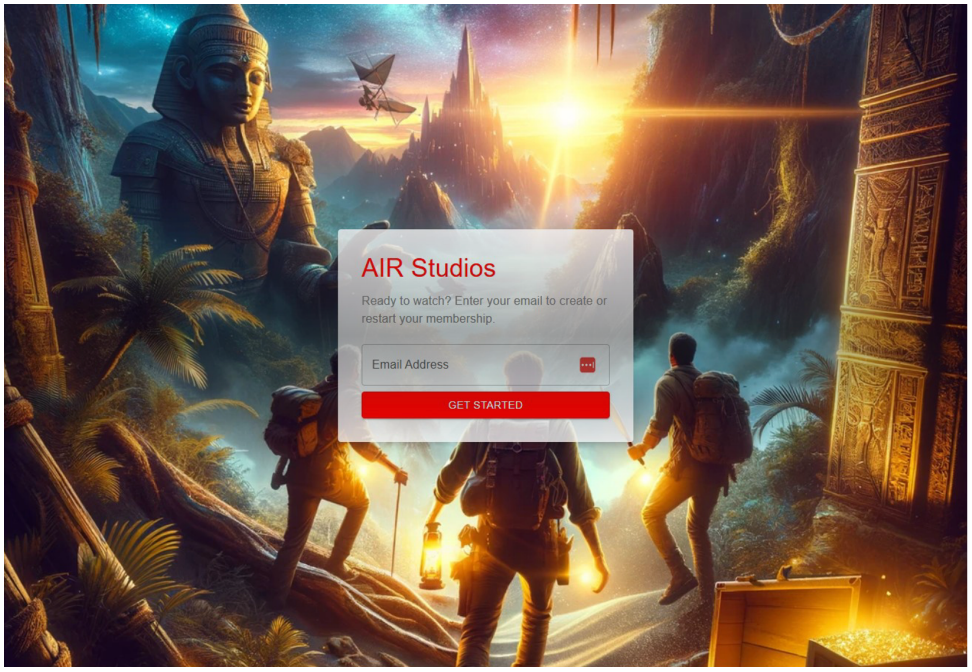

After the email address and password prompts, the website presented different content to participants based on their different experimental conditions. In the control condition, participants were asked to choose between two plans: "Standard with Ad Breaks" and "Standard with Personalized Ad Breaks." The two plans listed the same monthly price, number of included videos, and number of supported devices. See Figure 2. Users who chose non-personalized ad breaks were nagged in the dark patterns nagging condition to reconsider, with a pop-up prompt asking: "Do you want to opt out of personalized ads?" with a white text on red background button saying "Confirm" and a red text on white background button saying "Cancel."

Figure 2: Initial Personalization Choice

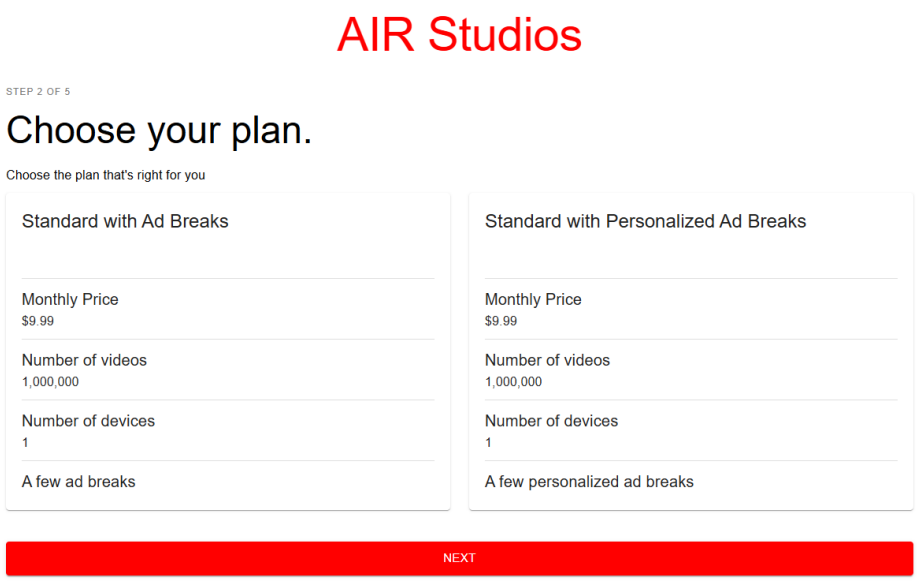

The next page asked participants in the control condition to choose their privacy settings and presented them with three toggles. One that turned on "strictly necessary cookies" was set to on, grayed out, and could not be adjusted. See Figure 3. This design mimics what consumers regularly encounter in consent management platforms. Another toggle that turned on "performance cookies" was set to on, but consumers could turn it off. A third toggle, "targeting cookies," was set to off and was adjustable. Participants could click a button labeled "Next" when they were ready to advance to the next screen.

Figure 3: Cookie Settings (Control on Top, Dark Patterns Below)

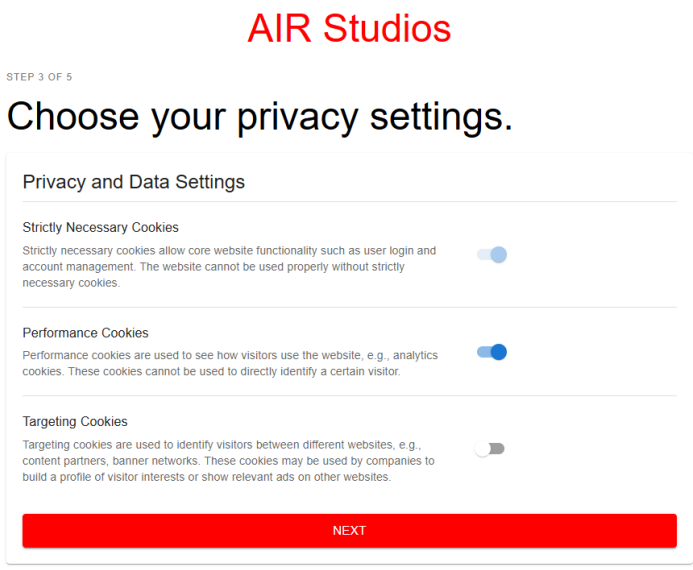

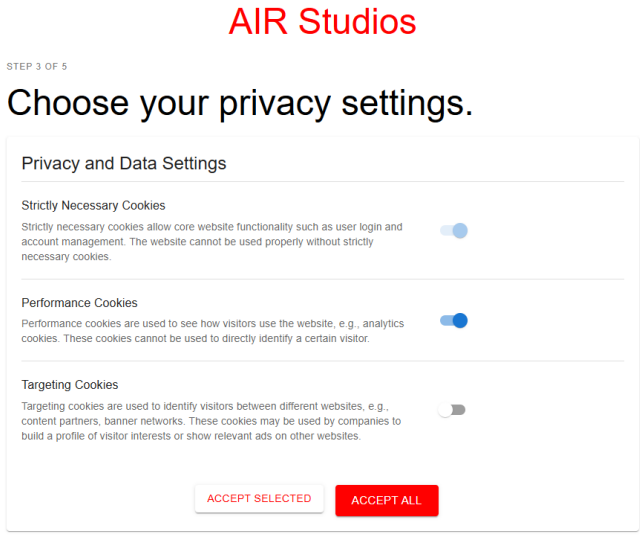

Participants then completed two screens asking filler questions, such as their zip code and genre preferences. Finally, the website said, "Congratulations on signing up! Click below to go to Settings." See Figure 4. The setting screen featured two grayed-out options—"Membership" and "Billing"—and one option that they were directed to click: "Privacy and Data Settings." Choosing the "Privacy and Data Settings" button took participants to the final screen, which contained three more toggles, all set to off. The toggles were labeled "allow matched identifier communications," "allow behavioral advertising," and "do not sell or share my personal information."[99]

Figure 4: Successful Sign-up Confirmation

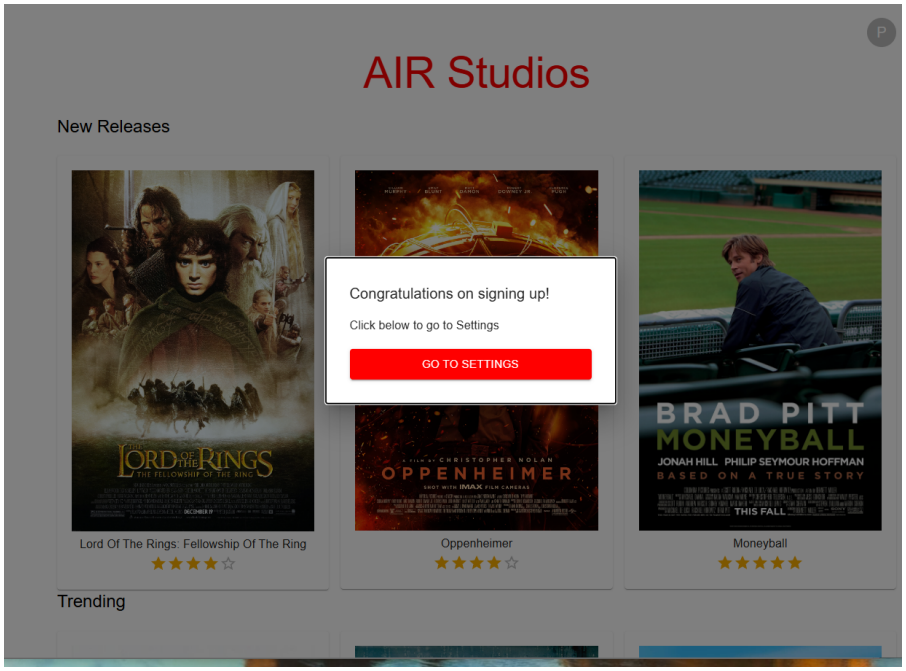

[99] The matched identifier option described itself as "Use privacy protected contact information from my AIR account to display relevant marketing on third party services." The behavioral advertising option described itself as "Behaviorally targeted advertising are ads selected based on your use and/or interactions with unaffiliated third-party websites and apps over time. This is tracked using your AIR Studios contact info and/or devices. You can choose not to receive behaviorally targeted ads on AIR Studios." The Do Not Sell option was left undefined.

These various options within the sign-up process generated six key choices: whether ads were standard or personalized, whether each of performance cookies, targeting cookies, matched identifiers, and behavioral advertising were on, and whether participants activated the Do Not Sell option.

There were two dark patterns conditions. These conditions presented the same choices, but did so in a way intended to guide participants to privacy-exposing outcomes. In the primary dark patterns condition, the following changes were made:

1. If participants selected the "Standard with Ad Breaks" plan, they were asked, "Do you want to opt out of personalized ads?" If they clicked the "Confirm" button, which was highlighted in red, they proceeded. If not, they were returned to make the choice again. This interface employed a nagging dark pattern (asking subjects who opted out of personalized ads whether they were certain, but not posing this question to subjects who preferred personalized ads). Dark pattern participants were nagged again about this choice after the genre selection screen. If they had not yet turned on personalized ads they were asked "Are you sure that you don't want personalized ad breaks?" with the options being "Personalize and continue" and "Maybe later." In addition to being a further nag, this prompt also incorporated interface interference by making the personalized ad breaks option more visually prominent. See Appendix 2.1 for images of these dark patterns.
2. On the cookie selection screen, the same options—strictly necessary and performance cookies—were toggled on, but the text on the button to advance to the next screen had changed. Rather than saying "Next," the options were "Accept All," which was highlighted in red, and "Accept Selected." Clicking the "Accept All" button also turned on targeting cookies. This interface employed an interface interference dark pattern (making the "Accept All" button more visually prominent) as well as a preselection dark pattern (making performance cookies on by default). See Figure 3. If users clicked on the "Accept Selected" button in the dark patterns condition, they could choose any combination of performance and targeting cookies

setting. If they selected "Accept All" performance cookies and targeting cookies would be activated.

3. If participants did not turn on targeting cookies, clicking "Accept Selected" prompted a screen that said, "Are you sure you want to opt out of advertising cookies?" If they clicked "Confirm," which was highlighted in red, they proceeded. If they clicked the "Cancel" button, they were returned to make the choice again and targeting cookies was toggled on. This interface employed a nagging dark pattern (asking users who made the privacy-protective choice whether they were sure they wanted to do that, but posing no such prompt to users who opted for less privacy).[100] See Appendix 2.2.
4. On the final settings screen, participants in the control group were shown a series of toggles but participants in the dark patterns condition were not presented with toggles. Instead, they saw Figure 5 below, which gave default "yes" answers to the settings for matched identifiers and behavioral advertising and a default "no" answer to the Do Not Sell option. Participants could either "Accept and Finish," which was highlighted in red, or "Edit Preferences." Clicking "Edit Preferences" gave participants access to the same toggles that were present in the control condition, though matched identifiers and behavioral advertising were now toggled to on by default. This interface combined a Preselection dark pattern (fewer privacy-protective choices were the default) and an obstruction dark pattern (selecting privacy-protective settings required subjects to click through an additional screen compared to waiving their privacy rights) for matched identifiers and behavioral advertising. For the "Do Not Sell or Share" toggle it combined preselection and obstruction (for the aforementioned reasons) with a confusion dark pattern (a double negative prompt that increased the cognitive demands placed on subjects).

[100] In both instances where nagging was employed (bullet points 1 and 3 above), the effect of the nag was potentially undermined by interface interference that encouraged them to confirm their earlier choice by making that button more visually prominent.

Figure 5: The Privacy and Data Settings (Control on top, Dark Patterns below)

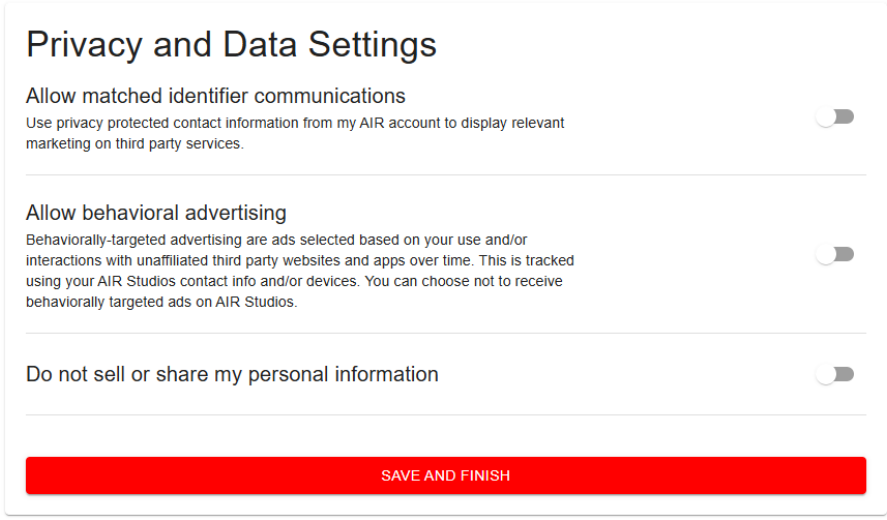

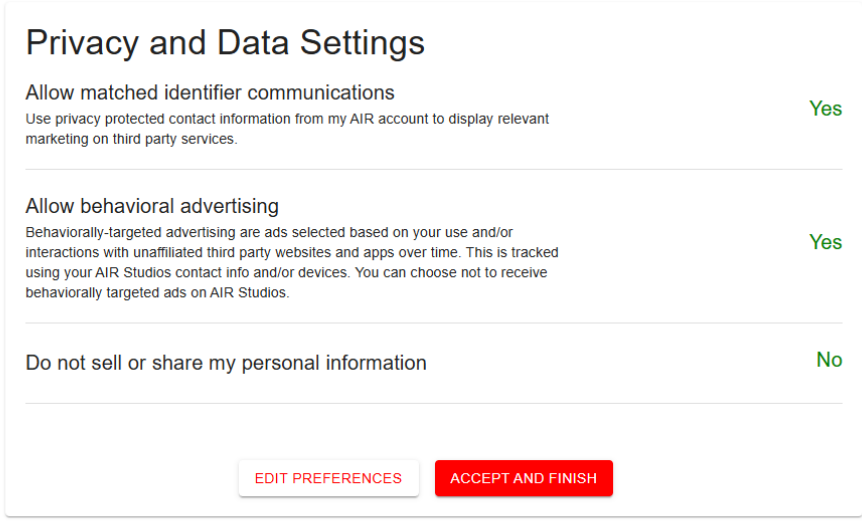

The alternative dark patterns condition removed the nagging dark patterns in items 1 and 3 above but kept the dark patterns in 2 and 4.

Participants in all three conditions (control and both dark patterns conditions) were then redirected to Qualtrics to finish a second part of the survey. This brief part

of the survey asked participants how difficult they found the sign-up process and how difficult they found choosing privacy protective options in the process. They then completed a mood measure, described below, and self-reported their interest in subscribing to the website based on what they had seen. Figure 6 displays a flow-chart of the different screens that experimental subjects saw, along with the dark patterns to which those in the experimental conditions were exposed.

Figure 6: Website Flow

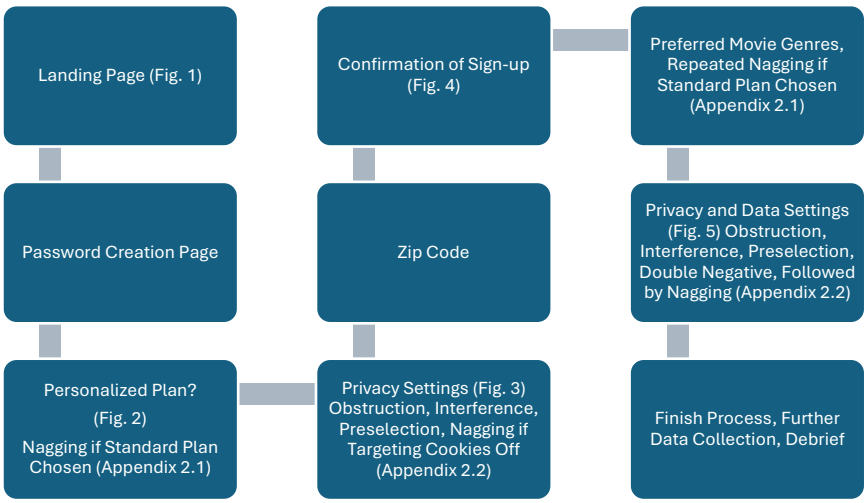

B. Participants

Participants began in one Qualtrics survey, went to a separate website to do the sign-up process, and then returned to another Qualtrics survey. This created three separate data files that were linked via a participant ID number, which was passed along as the participant proceeded through the task.[101]

[101] Before assessing any of the attention checks, the data needed to be processed to link participants across files and to deal with any duplicate cases. A small number of participants started the first survey ("Part I") more than once. Any participant who appeared in Part I more than once was evaluated. Their data was discarded if they reached the condition assignment more than once and were assigned to different goal conditions (as they would have seen the other goal prompt); if their data was inconsistent in the demographics section (no one fell into this category, but it was checked); or if they actually entered the test website more than once (as this would have made it unclear which data to use in analysis). This removed very few people as the most common

On most measures, data quality appeared high. Consistent with our registration, participants were removed for progressing through the study too quickly, which was defined as in less than one-third the median time; reporting inconsistent answers on a pair of questions several screens apart asking about how many cats or dogs lived in their household; or for giving a gibberish response when asked for comments or suggestions on the website at the close of the study. These checks excluded three speeders, nine people who gave irreconcilable responses to the pet ownership question, and zero people giving gibberish comments. These attention checks reduced our sample from 1,743 subjects to 1,731. Another 9.9% of these 1,731 subjects misreported their assigned goal (privacy protective responses or normal behavior) and were therefore excluded.[102] This left a final sample of 1,560 subjects.

We sought to recruit a census-representative sample in terms of age, gender, race and ethnicity, and educational attainment. As illustrated in Appendix I, we were successful in terms of age, gender, and race and ethnicity. We fell a little short in the category with the least educational attainment (people who neither completed high school nor earned a G.E.D.), however. Despite not achieving a perfectly representative distribution, we still had 28.1% of the sample with either a high school diploma or less as their highest completed level of education. We were also only slightly over-representing the most educated subset—those with a graduate-level degree (15.1% achieved versus 14.2% targeted).

C. Goals and the Effectiveness of Dark Patterns

Our primary study questions were whether dark patterns were effective in influencing choices consumers make when navigating through their privacy settings

person in this category was one who started Part I, did virtually none of it, and then restarted it a few minutes later.

[102] Some of these research subjects may have been excluded because of quite poor reading comprehension skills or low cognitive ability. We hypothesize that these individuals may have been especially susceptible to dark patterns were they included in the sample. Accordingly, our attention check likely caused our reported results to underestimate the potency of dark patterns on the population writ large.

and whether that effectiveness was limited by privacy goals. The overall results are illustrated in Table 1.

Table 1: Percentage of participants opting for each of the below features by condition.

(Higher Percentages indicate more of the subjects waiving privacy protections, except for the "Do Not Sell" Interface, where lower percentages indicate more of the subjects waiving privacy protections)

Interface (Relevant dark patterns in parentheses)	**Goal**	**Control**	**Dark patterns**	**Dark patterns, but no nagging**
Personalized Ads (Nagging only)	Normal	51.8%	64.9%	52.4%
	Protect Privacy	29.7%	51.5%	30.2%
	Overall	40.3%	58.0%	40.5%
Performance Cookies (Interface interference + Preselection)	Normal	54.9%	63.0%	68.8%
	Protect Privacy	27.1%	41.6%	41.9%
	Overall	40.5%	52.1%	54.4%
Targeting Cookies (Interface interference + Nagging)	Normal	9.0%	36.3%	38.1%
	Protect Privacy	7.0%	27.4%	26.4%
	Overall	8.0%	31.7%	31.9%
Matched Identifiers (Preselection + Obstruction + Interface interference)	Normal	8.2%	51.9%	61.9%
	Protect Privacy	7.3%	24.5%	24.5%
	Overall	7.8%	37.9%	41.9%
Behavioral Targeting (Preselection + Obstruction + Interface interference)	Normal	11.8%	51.1%	60.2%
	Protect Privacy	5.5%	21.2%	20.8%
	Overall	8.5%	35.8%	39.1%
Do Not Sell (Preselection + Obstruction + Confusion)	Normal	71.4%	40.5%	43.3%
	Protect Privacy	81.3%	62.4%	57.7%
	Overall	76.5%	51.7%	51.0%

Note: The only dark pattern targeting personalized ads was a nagging pattern, so there was no dark pattern targeting that measure in the "no nagging" condition.

We assessed the significance of the cross-condition differences using a series of binary logistic regressions. These use goal condition, dark patterns condition, and their interaction as predictors of each of the six dependent measures, which are binary (on or off). For simplicity, the two dark patterns conditions will be combined for all measures except personalized ads, for which the no nagging dark patterns condition was identical to the control condition from the standpoint of the participants. Whether the dark patterns included nagging had no effect on any of the other five measures.[103] A binary logistic regression effectively tests whether the chance of a particular outcome (for example, a privacy setting being on) varies depending on some set of predictor variables. Thus, this test shows that whether a person was in the dark patterns condition or the dark patterns without nagging condition did not reliably alter the odds that they would turn any of the other settings on.

Many of these analyses look at the effect of the goal, the effect of dark patterns conditions, and their statistical interaction. A statistical interaction tests whether the effect of one factor in the analysis depends upon the value of the other. Imagine a study that examines the effect of adding more sugar and more salt to a cookie recipe. When there is a low level of salt, people like it when there is more sugar as opposed to less. When there is a high level of salt, the cookie is terrible regardless of how much sugar is added. So, adding more sugar only helps when the cookie is not heavily oversalted.

Here, we will be asking whether dark patterns work as well when people have been given a privacy goal. This involves looking at the effects of dark patterns, the effects of privacy goals, and the effect of their interaction — whether the effect of being exposed to dark patterns changes depending on whether someone has a privacy-maximizing goal.

---

[103] This was assessed using a binary logistic regression that compared dark patterns with nagging with dark patterns without nagging (1, 0). This factor was not statistically significant on any of the 5 measures.
Performance cookies B = -0.10, Wald = 0.59, $p = 0.443$, Exp(B) = 0.91, 95% CI [0.71, 1.16]
Targeting cookies B = -0.01, Wald = 0, $p = 0.962$, Exp(B) = 0.99, 95% CI [0.76, 1.29]
Matched identifier B = -0.17, Wald = 1.77, $p = 0.183$, Exp(B) = 0.84, 95% CI [0.66, 1.08]
Behavioral advertising B = -0.14, Wald = 1.19, $p = 0.275$, Exp(B) = 0.87, 95% CI [0.68, 1.12]
Do not sell information B = 0.03, Wald = 0.05, $p = 0.829$, Exp(B) = 1.03, 95% CI [0.8, 1.31].

1. Dark Patterns and Goals Affected Ad Personalization Choices

Personalized ads. There was a significant effect of goal condition, a significant effect of dark patterns condition, and a nonsignificant trend toward interaction between the two.[104] This means that participants with a privacy goal were significantly less likely to opt for personalized ads and participants exposed to the nagging dark pattern were more likely to opt for them. Dark patterns actually had a larger effect here when participants had a privacy goal, though this was a nonsignificant trend ($p = .086$).[105]

Most notable on this measure is the high base rate in the normal goal default condition. About half the participants, left entirely to themselves, opted for personalized ads. This was made slightly more common by the use of dark patterns. Giving people an explicit privacy goal caused fewer people to select personalized ads when dark patterns were not present, but dark patterns substantially increased uptake.

2. Dark Patterns and Goals Affected Cookie Selections

Performance cookies. This setting was on by default in all conditions. This would be off if participants turned it off and, in the dark patterns condition, also chose the "Accept Selected" button rather than the "Accept All" button on the cookies page. Thus, subjects were exposed to preselection and interface interference dark patterns. Combining the two dark patterns conditions, there was a significant effect of dark patterns versus not and a significant effect of goal, but no interaction.[106] Dark patterns increased the rate of selecting performance cookies by 12.7 percentage

---

[104] Goal: B = 0.93, Wald = 50.9, $p < 0.001$, Exp(B) = 2.54, 95% CI [1.97, 3.29]
Dark Pattern: B = 0.91, Wald = 35.22, $p < 0.001$, Exp(B) = 2.48, 95% CI [1.84, 3.35]
Interaction: B = -0.38, Wald = 2.94, $p = 0.086$, Exp(B) = 0.69, 95% CI [0.45, 1.06].

[105] This is likely due to the low percentage of people opting for personalized ads in the privacy goal-no dark patterns condition. Because that number was so low, there was a lot of room for dark patterns to work. With about half the sample already opting for personalized ads in the control-goal condition, dark patterns had less space for improvement.

[106] Goal: B = 0.98, Wald = 58.23, $p < 0.001$, Exp(B) = 2.68, 95% CI [2.08, 3.44]
Dark Pattern: B = -0.66, Wald = 16.44, $p < 0.001$, Exp(B) = 0.52, 95% CI [0.38, 0.71]
Interaction: B = 0.2, Wald = 0.8, $p = 0.372$, Exp(B) = 1.22, 95% CI [0.79, 1.9].

points and a privacy goal reduced the rate of selecting performance cookies by 25.4 percentage points.

Targeting cookies. This was off by default in the control condition. It would be on if participants turned it on, in any condition, or if they clicked "Accept All" in the dark patterns conditions. Thus, in the dark patterns conditions subjects had to overcome interface interference and nagging dark patterns. Again, there was a significant effect of dark patterns (both conditions combined) versus not and a significant effect of goal, but no interaction.[107] Dark patterns increased the acceptance of targeting cookies by 23.8 percentage points and a privacy goal reduced the acceptance rate for targeting cookies by 7.5 percentage points.

Looking at the base rate differences between performance and targeting cookies shows the power of defaults. Performance cookies were on by default and are much more common across all conditions. Targeting cookies are off by default and were rarely turned on, except when dark patterns are present.

3. Dark Patterns and Goals Affected Privacy Settings, with Significant Interaction Effects

Matched identifiers. This setting was off in the control condition and on in dark patterns, creating a preselection dark pattern, and dark patterns also obstructed efforts to edit this feature. Here, the results differed. There was a significant effect of dark patterns, a significant effect of goal, and a significant interaction between dark patterns and subject goals.[108] In the normal goal condition, matched identifiers were on 8.2% of the time in the control condition and 56.9% of the time in the two dark patterns conditions (a 48.7 percentage point difference). In the privacy protective condition, matched identifiers were on 7.3% of the time in control and only 24.5% of the time in the dark patterns conditions (a 17.2 percentage point difference). Overall, dark patterns had a smaller effect when there was a privacy

---

[107] Goal: B = 0.47, Wald = 12.32, $p < 0.001$, Exp(B) = 1.6, 95% CI [1.23, 2.09]
Dark Pattern: B = -1.59, Wald = 38.46, $p < 0.001$, Exp(B) = 0.2, 95% CI [0.12, 0.34]
Interaction: B = -0.19, Wald = 0.3, $p = 0.585$, Exp(B) = 0.83, 95% CI [0.42, 1.64].

[108] Goal: B = 1.39, Wald = 105.83, $p < 0.001$, Exp(B) = 4.02, 95% CI [3.08, 5.24]
Dark Pattern: B = -1.41, Wald = 31.14, $p < 0.001$, Exp(B) = 0.24, 95% CI [0.15, 0.4]
Interaction: B = -1.26, Wald = 12.88, $p < 0.001$, Exp(B) = 0.28, 95% CI [0.14, 0.56].

goal. Nevertheless, the effect of dark patterns was still significant in the privacy protective goal condition.[109]

Behavioral targeting. This measure had the same defaults and dark patterns as matched identifiers. Again, there was a significant effect of dark patterns, a significant effect of goal, and a significant interaction. [110] In the normal goal condition, behavioral targeting was on 11.8% of the time in the control and 55.7% of the time in the two dark patterns conditions (a 43.9 percentage point difference). In the privacy protective condition, behavioral targeting was on 5.5% of the time in control and 21.0% of the time in the dark patterns conditions (a 15.5 percentage point difference). So dark patterns had a smaller effect when there was a privacy goal. Nevertheless, the effect of dark patterns was still significant and quantitatively substantial in the privacy protective goal condition.[111]

We applied our most extreme dark patterns to the matched identifiers and behavioral targeting items by changing to an anti-privacy default, obstructing an alteration of that default, and using visual salience to discourage editing. In the normal goal condition, this was highly effective in causing people to be opted in to matched identifiers and behavioral targeting. This effectiveness was more than cut in half by a privacy goal, however. The percentage point change was 2.83 times greater when a privacy goal was absent.

Do not sell or share information. This setting was turned off by default in all conditions, and dark patterns also obstructed efforts to edit this feature. There was a significant effect of dark patterns and a significant effect of goal, but no interaction.[112] Dark patterns decreased the use of Do Not Sell by 25.2 percentage points and a privacy goal increased the use of Do Not Sell by 15.4 percentage points.

---

[109] χ^2 (1, N = 812) = 35.09, p < .001.

[110] Goal: B = 1.54, Wald = 122.67, p < 0.001, Exp(B) = 4.68, 95% CI [3.56, 6.15]
Dark Pattern: B = -1.52, Wald = 28.19, p < 0.001, Exp(B) = 0.22, 95% CI [0.13, 0.38]
Interaction: B = -0.71, Wald = 3.98, p < 0.05, Exp(B) = 0.49, 95% CI [0.24, 0.99].

[111] χ^2 (1, N = 812) = 32.66, p < .001.

[112] Goal: B = 0.47, Wald = 12.32, p < 0.001, Exp(B) = 1.6, 95% CI [1.23, 2.09]
Dark Pattern: B = -1.59, Wald = 38.46, p < 0.001, Exp(B) = 0.2, 95% CI [0.12, 0.34]
Interaction: B = -0.19, Wald = 0.3, p = 0.585, Exp(B) = 0.83, 95% CI [0.42, 1.64].

The Do Not Sell measure was unusual in that it was the only feature that protected privacy when it was turned on as opposed to off. This linguistic change, and the double negative used in the dark patterns condition, created the potential for consumers to become confused about what choice they were making. Even though the Do Not Sell option was off by default in all conditions, it was still wildly popular. More than 71% of people in the control condition turned it on even when they were told to simply choose whatever options they normally would. (More than 81% of users in the privacy protective condition turned it on.) Our data thus identifies an important exception to the generally sticky nature of default settings.[113]

The enormous popularity of the "Do not sell or share my personal information" option is a particularly policy-relevant finding because one of the CCPA's main policy provisions is a requirement that websites subject to the law provide consumers with a readily accessible link or toggle that allows them to prevent companies from selling or sharing their personal information.[114] Our results provide some indication of why this popular consumer option is one that websites and apps might go to great lengths to try to thwart through the use of dark patterns. It's notable here that this apparently strong privacy preference, which was made stronger if people were given a privacy goal, was substantially frustrated by dark patterns that made the option more difficult to edit and more confusing.

Summarizing the results in Table 1, the effectiveness of the dark patterns tested jumps off the page. Relatively subtle differences in the visual appearance of user interfaces sometimes had the effect of tripling or even quadrupling the overall percentage of subjects who selected less privacy protective options, and even when subjects had instructions to maximize their privacy protections, the dark patterns often tripled the percentage of people who waived their privacy rights. Some of the largest effects arose from dark patterns arising outside of the cookie/CMP context, which are the type of interface most frequently studied by previous dark patterns researchers. In terms of social science research findings, these effect sizes from interface interference, obstruction, and preselection dark patterns are enormous.

---

[113] *See generally* Ian Ayres, *Regulating Opt-Out: An Economic Theory of Altering Rules*, 121 YALE L. J. 2032 (2012).

[114] Cal. Civ. Code § 1798.135(a)(1) (West 2025).

It's no wonder that these dark patterns have sparked so much alarm among regulators and legislators.

Overall, there were few effects on the post-website survey measures. Participants completed a short-form of the Positive and Negative Affect Schedule (PANAS), but there were no significant cross-condition differences on positive or negative mood. There were also no significant differences on interest in subscribing to the website. Having a privacy goal made participants believe that signing up for the website and choosing the most privacy protective options were both harder.[115] Also, the dark patterns with nagging condition was perceived as more difficult than the other two on both measures, though dark patterns without nagging did not differ from the control.[116] There were no other differences and no interactions. These dependent variables are important findings in and of themselves because they indicate that the use of numerous dark patterns did not generate a meaningful backlash from potential customers. This finding replicates Luguri & Strahilevitz's determination that there is not a meaningful penalty in the marketplace for firms that employ relatively subtle dark patterns.[117]

D. Behavior in the Dark Patterns Conditions

The website logged every click participants made throughout the study, which allowed us to see how participants responded to each biased choice presented to

---

[115] This was assessed using a 3 (Dark patterns condition) by 2 (Goal condition) between-subjects ANOVA. ANOVAs are statistical tests that allow for the comparison of means across experimental conditions. Here the test is asking whether people in the privacy protective goal condition have different scores on each of these measures compared to people in the normal goal condition. The next footnote looks at the comparison across dark patterns conditions.
Effect of goal on:
Sign up difficulty $F(1,1553) = 14.72$, $p < 001$, $\eta^2 = 0.009$, Normal ($M = 1.25$, $SD = 0.5$); Privacy protective ($M = 1.36$, $SD = 0.63$).
Privacy choices difficulty $F(1,1553) = 10.06$, $p = 0.002$, $\eta^2 = 0.006$, Normal ($M = 1.53$, $SD = 0.77$); Privacy protective ($M = 1.67$, $SD = 0.88$).

[116] Sign up difficulty $F(2,1553) = 4.84$, $p = 0.008$, $\eta^2 = 0.006$. Control ($M = 1.30$, $SD = 0.59$); Dark patterns without nagging ($M = 1.26$, $SD = 0.48$); Dark patterns with nagging ($M = 1.37$, $SD = 0.64$).
Privacy choices difficulty $F(2,1553) = 11.33$, $p < .001$, $\eta^2 = 0.014$. Control ($M = 1.5$, $SD = 0.77$); Dark patterns without nagging ($M = 1.57$, $SD = 0.76$); Dark patterns with nagging ($M = 1.73$, $SD = 0.93$).

[117] *See supra* text accompanying note 59.

them in the dark patterns conditions. Recall that on the cookies page, dark patterns participants were given the choice between "Accept All" and "Accept Selected." More participants "accepted all" in the normal goal condition (35.3%) than in the privacy protective goal condition (26.3%).[118] Similarly, more participants in the normal goal condition chose "Save and Finish" rather than "Edit" when confronted with less privacy-protective choices on the final settings page (41.2% in normal goal versus 13.0% in privacy protective goal).[119]

The behaviors in response to the nagging dark patterns were more complex. Recall that there were two nags related to the choice to have personalized ads. In response to the first nag, 97.45% of participants persisted in rejecting personalized ads. In response to the second nag, which appeared several screens later, 30.3% of participants in the normal goal condition and 21.8% of participants in the privacy goal condition changed to opting for personalized ads (25.5% overall). These proportions did not significantly differ.[120] Finally, 95% of participants who received the nag asking if they were sure they did not want targeting cookies rejected that nag. These varied findings on the efficacy of nagging are important, novel, and (at first blush) somewhat puzzling. Our working hypothesis is that nagging dark patterns work best when the nag is combined with another dark pattern, such as interface interference or preselection. Because the present study does not allow us to isolate precisely when nagging is most effective, our research team is in the process of running follow-up experiments to resolve those questions. In our initial data, we find that people who are directed to maximize their privacy protections have an easier time overcoming the effects of nagging dark patterns than they do overcoming less blatant dark patterns like preselected defaults and interface interference.

E. Limited Role of Individual Differences

Prior experimental work has sometimes shown that individual differences can moderate the effect of dark patterns. For instance, Luguri and Strahilevitz showed that people with lower educational attainment are more likely to be influenced by

---

118 χ^2 (1, N = 1032) = 9.71, p = .002.

119 χ^2 (1, N = 1032) = 105.17, p < .001.

120 χ^2 (1, N = 302) = 2.75, p = .091.

dark patterns, though Zac et al. found only weak effects of educational attainment.[121] Moreover, two relatively small-*n* research studies showed that older people are more likely to be influenced by dark patterns.[122] These effects are not always present, however. For instance, neither team of researchers reproduced the effect shown by the other, so it appears to be the case that the role of individual differences varies across contexts.

Here, we measured a broad array of individual differences including educational attainment, age, self-reported social class, political orientation on a liberal-to-conservative axis, self-assessed technology skills, and right-wing authoritarianism.

Prior work has observed that attitudes about governmental searches are correlated with the social psychological construct known as right-wing authoritarianism, with authoritarians being less privacy protective.[123] The social psychological theory of authoritarianism defines authoritarians as people who are especially willing to submit to authority; who believe it is particularly important to yield to traditional conventions and norms; and who are hostile and punitive toward those who question authority or who violate traditional conventions and norms.[124] The specific authoritarianism scale used in prior work,[125] and again employed here, is the Authoritarian Submission scale. This scale is intended to measure the first of the defined impulses: the extent to which people believe authority should be

---

121 *Compare* Luguri & Strahilevitz, *supra* note 2, at 70; *with* Zac et al., *supra* note 67, at 19–23.

122 Anaraky et al., *supra* note 92, at 11; Woon Chee Koh & Yuan Zhi Seah, *Unintended Consumption: The Effects of Four E-Commerce Dark Patterns*, 11 CLEANER & RESPONS. CONSUMP. 1001 (2023).

123 *See*, e.g., Matthew B. Kugler & Lior J. Strahilevitz, *Actual Expectations of Privacy, Fourth Amendment Doctrine, and the Mosaic Theory*, 2015 SUP. CT. REV. 205, 252–55 (2015); Matthew B. Kugler & Mariana Oliver, *Constitutional Pandemic Surveillance*, 111 J. CRIM. L. & CRIMINOLOGY 909, 938–40 (2021); Matthew B. Kugler, *Public Perceptions Can Guide Regulation of Public Facial Recognition*, 25 COLUM. SCI. & TECH. L. REV. 1, 35–36 (2023).

124 *See* Bob Altemeyer, *The Other "Authoritarian Personality"*, 30 ADVANCES IN EXPERIMENTAL SOC. PSYCH. 47 (Mark P. Zanna ed., 1998).

125 *See* Kugler & Strahilevitz, *supra* note 123; Kugler & Oliver, *supra* note 123.

respected and obeyed rather than challenged and questioned.[126] Here, the scale exhibited high reliability.[127]

The technology skills measure was adapted from a measure used by Isabel Rodríguez-de-Dios and colleagues and revised to target the kinds of technology skills most relevant to web navigation.[128] Here, the scale exhibited acceptable reliability.[129]

We conducted a series of binary logistic regressions that looked at the main effects of each of these individual difference measures and their interactions with dark patterns on each of the main dependent measures.[130] As illustrated in Table 2, individual differences played a relatively modest role. The only consistent main effects were on authoritarianism (five of the six measures) and self-reported

---

[126] We measured this at the beginning of the survey, prior to the website task. Scale items were presented in random order and included "It's great that many young people today are prepared to defy authority" (reverse coded) and "What our country needs most is discipline, with everyone following our leaders in unity." The response scale ranged from 1 (strongly disagree) to 5 (strongly agree). Higher scores indicate stronger endorsement of authoritarian ideologies. John Duckitt et al., *A Tripartite Approach to Right-Wing Authoritarianism: The Authoritarianism-Conservatism-Traditionalism Model*, 31 POL. PSYCH. 685, 690 (2010) ("Thus, the 'authoritarian submission' dimension can be defined as expressing attitudes favouring uncritical, respectful, obedient, submissive support for existing societal or group authorities and institutions (protrait) versus critical, questioning, rebellious, oppositional attitudes to them (contrait).").

[127] Cronbach's alpha = .867.

[128] Isabel Rodríguez-de-Dios et al., *Development and Validation of a Digital Literacy Scale for Teenagers*, 4 INT'L CONF. ON TECH. ECOSYSTEMS FOR ENHANCING MULTICULTURALITY 1067 (2016). A version of this scale was previously used by Zac et al., *supra* note 67, at 14–15, 37.
The response scale ranged from 1 (strongly disagree) to 5 (strongly agree). Higher scores indicate stronger self-reported technology skills.
Protrait items were:

- I know how to bookmark a website so I can view it later.
- I know how to mark unwanted emails as spam.
- I usually know how to change privacy settings on websites.

Contrait items were:

- Sometimes I end up on websites without knowing how I got there.
- I find the design of many websites to be confusing.
- I often ask people for help when I need to install new applications on my computer or phone.

[129] Cronbach's alpha = .664.

[130] This analysis used effects coding for the dark patterns measure (-1 for no patterns and +1 for patterns) and z scored the individual difference measures to center and standardize them.

technology skills (four of the six measures), with authoritarianism associated with anti-privacy responses and technology skills associated with pro-privacy responses. In other words, people who exhibit authoritarian personality traits are more vulnerable to being manipulated by dark patterns, and people who report that they are relatively competent at working with technologies are less vulnerable.[131]

Interactions with the dark patterns conditions were rare, and our experiment provides no consistent evidence of certain kinds of people being especially vulnerable to dark patterns in this study. Educational attainment interacted with dark patterns only on the matched identifier measure. Participants lower in educational attainment were more likely to have matched identifiers turned on in the dark patterns condition, but there was no effect of educational attainment in the control condition.[132] The pattern was similar for technology skills and the Do Not Sell measure. Those higher in technology skills were more likely to have Do Not Sell turned on in the dark patterns condition, but there was no effect in the control condition.[133]

Authoritarianism interacted with dark patterns only on the targeting cookies measure. Here, authoritarianism had a larger effect in the control condition than in the dark patterns condition, though it was significant in both.[134] People higher in authoritarianism were more likely to turn on this feature, especially in the baseline condition, where fewer people did so overall. In short, and contrary to some findings

---

[131] CloudResearch is an online research platform so we suspect that the least technologically literate American adults were not represented in our pool of subjects, even though we have a census-weighted across the dimensions that we can measure. This problem is not unique to our study, of course. In any event, it is not obvious that the responses of Americans who are rarely online or lack the resources to afford internet access to dark patterns are especially policy-relevant. We are studying the behavioral responses of the kinds of people who do get exposed to dark patterns in the real world, not those who don't and won't receive such exposure.

[132] Educational attainment on matched identifiers:
Control B = 0.21, Wald = 1.53, $p = 0.217$, Exp(B) = 1.23, 95% CI [0.89, 1.71]
Dark patterns B = -0.13, Wald = 4.37, $p < 0.05$, Exp(B) = 0.88, 95% CI [0.77, 0.99].

[133] Technology skills on Do Not Sell:
Control B = -0.09, Wald = 0.79, $p = 0.373$, Exp(B) = 0.91, 95% CI [0.74, 1.12]
Dark patterns B = 0.21, Wald = 10.80, $p = 0.001$, Exp(B) = 1.24, 95% CI [1.09, 1.40].

[134] Authoritarianism on targeting cookies:
Control B = 0.53, Wald = 10.15, $p < 0.01$, Exp(B) = 1.71, 95% CI [1.23, 2.37]
Dark patterns B = 0.33, Wald = 24.52, $p < 0.001$, Exp(B) = 1.4, 95% CI [1.22, 1.59].

elsewhere in the dark-patterns literature,[135] the similarities among people in terms of vulnerability to dark patterns seem greater than the differences. Dark patterns succeed in tricking or manipulating people across the board, at least in our dataset.

Table 2: Individual differences and their interactions with dark patterns.

	Personalized Ads		Performance Cookies		Targeting Cookies	
	Exp(B)	CI	Exp(B)	CI	Exp(B)	CI
Dark Patterns (DP)	1.47***	[1.32, 1.65]	1.30***	[1.17, 1.46]	2.47***	[2.04, 3.00]
Authoritarianism	1.45***	[1.27, 1.65]	1.54***	[1.34, 1.77]	1.78***	[1.43, 2.21]
Tech Skills	0.86*	[0.77, 0.97]	0.76***	[0.68, 0.86]	0.91	[0.77, 1.07]
Age	0.95	[0.85, 1.06]	1.13*	[1.01, 1.27]	0.86	[0.71, 1.03]
Education	1.02	[0.91, 1.16]	0.85*	[0.75, 0.97]	0.98	[0.80, 1.19]
Political Orientation	0.99	[0.87, 1.13]	0.88	[0.77, 1.00]	0.76*	[0.61, 0.94]
Social Class	1.04	[0.92, 1.18]	0.91	[0.80, 1.03]	1.14	[0.94, 1.39]
DP by Authoritarianism	0.97	[0.85, 1.1]	0.91	[0.79, 1.05]	0.80*	[0.64, 0.99]
DP by Tech Skills	0.98	[0.87, 1.1]	1.01	[0.89, 1.13]	0.91	[0.77, 1.07]
DP by Age	1.01	[0.90, 1.13]	0.98	[0.87, 1.10]	1.04	[0.86, 1.25]
DP by Education	1.02	[0.90, 1.16]	1.01	[0.89, 1.15]	0.97	[0.79, 1.18]
DP by Political Orientation	0.94	[0.83, 1.07]	1.00	[0.88, 1.15]	1.20	[0.97, 1.48]
DP by Social Class	1.02	[0.91, 1.16]	1.02	[0.90, 1.16]	0.93	[0.77, 1.13]

	Matched Identifier		Behavioral Advertising		Do Not Sell	
	Exp(B)	CI	Exp(B)	CI	Exp(B)	CI
Dark Patterns (DP)	3.11***	[2.53, 3.81]	2.65***	[2.20, 3.18]	0.57***	[0.51, 0.64]
Authoritarianism	1.72***	[1.38, 2.15]	1.68***	[1.36, 2.07]	0.89	[0.77, 1.03]
Tech Skills	0.73***	[0.63, 0.86]	0.78**	[0.67, 0.91]	1.05	[0.92, 1.19]
Age	0.80*	[0.66, 0.97]	0.88	[0.73, 1.05]	1.13*	[1.00, 1.29]
Education	1.09	[0.89, 1.34]	0.93	[0.77, 1.13]	0.97	[0.84, 1.11]
Political Orientation	0.97	[0.79, 1.20]	0.84	[0.68, 1.03]	0.99	[0.86, 1.13]
Social Class	1.01	[0.82, 1.23]	0.97	[0.80, 1.18]	0.92	[0.80, 1.05]
DP by Authoritarianism	0.85	[0.68, 1.07]	0.87	[0.71, 1.07]	0.93	[0.81, 1.08]
DP by Tech Skills	0.96	[0.82, 1.13]	0.91	[0.78, 1.06]	1.18*	[1.04, 1.33]
DP by Age	1.07	[0.88, 1.30]	0.91	[0.76, 1.09]	1.12	[0.99, 1.27]
DP by Education	0.81*	[0.66, 0.99]	1.01	[0.83, 1.22]	0.95	[0.83, 1.09]
DP by Political Orientation	0.97	[0.79, 1.19]	1.16	[0.95, 1.42]	1.04	[0.91, 1.20]
DP by Social Class	1.05	[0.86, 1.29]	1.09	[0.90, 1.32]	1.04	[0.91, 1.19]

Note: For personalized ads, dark patterns refers only to the dark patterns with nags condition. For all other variables, it refers to both dark patterns conditions. N for all analyses is 1560. *** means $p < .001$; ** means $p < .01$; * means $p < .05$.

[135] *See* Luguri & Strahilevitz, *supra* note 2; Zac, *supra* note 67; Anaraky, *supra* note 92; Koh & Seah, *supra* note 122.

F. Follow-up on Nagging

Due to the importance of nagging dark patterns as a policy matter, and the lack of a literature evaluating the effectiveness of nagging dark patterns, we wanted to further explore the effectiveness of nags. In particular, we sought to isolate the effects of nagging from those of other dark patterns and examine whether different types of nags were effective to different degrees. We therefore altered the design of the streaming website to allow for the creation of three new conditions. These, along with the control condition (which persisted from the main study) formed the instrument for the nagging follow-up.

The three new nagging conditions (see Appendix 3 for figures) were as follows:

1. Repeated nags that used interface interference to highlight the preferred choice. These nags were a) about the subscription, on the subscription page, b) about targeting cookies, on the targeting cookies page, and c) about the subscription, on the genre page.
2. The same repeated nags, but without interface interference.
3. A single nag about the subscription, on the genre page, that uses interface interference.

1. Procedure and Participants

Except as noted, the study proceeded as before. As in the main study, participants were recruited on Cloud Research and entered the study via Qualtrics. All participants were instructed to "choose **whatever options you think you normally would**" during the signup process. They then were passed to the website and made the same choices as before. What changed was the structure of the nags. These new conditions were all based upon the control condition from the initial study, so people did not need to overcome preselection or obstruction dark patterns.

In the control and single nag conditions, the initial choice of subscription (personalized ads or not) was left undisturbed. In the two repeated nags conditions, participants who chose non-personalized ads received an immediate popup "Are you sure you want to opt-out of personalized ads?" The choices were "Show me personalized ads" and "Don't show me personalized ads." The interface interference condition highlighted the "Show me personalized ads" option in red. If someone

clicked on the "Show me personalized ads" option, this automatically enrolled them in personalized ads.

The experience on the cookie page was similar. If, in the two conditions with repeated nags, the participant did not turn on "targeting cookies" they were asked "Do you really want to block advertising cookies?" The choices were "Allow advertising cookies" and "Block them for now and ask me later." The interface interference condition highlighted the "Allow advertising cookies" option in red.

Finally, on the genre page, participants in all three nagging conditions who had not previously turned on personalized ads were nagged. After selecting their preferred genres, participants were asked "Are you sure that you don't want personalized ad breaks?" The options were "Personalize and continue" and "Maybe later." In the two interface interference conditions, the "Personalize and continue" option was highlighted in red.

The remaining screens continued as before and did not include any nags.

Following the example of the first study, participants were removed for progressing through the study too quickly, which was defined as in less than one-third the median time; reporting inconsistent answers on a pair of questions several screens apart asking about how many cats or dogs lived in their household; or for giving a gibberish response when asked for comments or suggestions on the website at the close of the study. These checks excluded two speeders, fourteen people who gave irreconcilable responses to the pet ownership question, and zero people giving gibberish comments. This reduced the sample from 929 to 913. Another 17.2% of this remaining sample misreported their assigned goal (normal behavior) and were therefore excluded. That left a final sample of 756 subjects. Full sample demographics are available in Appendix 1.

2. The Effects of Different Kinds of Nags

Overall, the nags resulted in greater adoption of both personalized ads and targeting cookies.[136] About seventeen percent of those nagged changed to accepting

[136] A chi square contrasting the three nagging conditions versus control (for personalized ads) and the two relevant nagging conditions vs. control and irrelevant nagging condition (for targeting cookies, recall that the single nag was later and only on something else), both show that nags

personalized ads. This percentage did not vary significantly based on nag type (see Table 3).[137] Interestingly, nagging appears to have diminishing returns. The single nag after the genre screen was as effective as the combination of the two nags in the other two conditions.

Table 3: Nagging Type Effectiveness – Personalized Ads

Accepted Personalized Ads

Condition	Initially	After 1st Nag	After 2nd Nag	Total	Improvement
Control	48.02%			48.02%	
Repeated Nags, No Interface Interference	52.28%	7.11%	1.52%	60.91%	8.63%
Repeated Nags, Interface Interference	46.45%	3.28%	3.83%	53.55%	7.10%
Single Nag, Interface Interference	42.71%		11.06%	53.77%	11.06%

Percent Nagged Who Changed

Condition	After 1st Nag	After 2nd Nag	Total[138]
Repeated Nags, No Interface Interference	14.89%	3.75%	18.09%
Repeated Nags, Interface Interference	6.12%	7.61%	13.27%
Single Nag, Interface Interference		19.30%	19.30%

As can be seen in Table 4, nags were similarly effective in changing responses to the targeting cookies setting. Here the effect was more pronounced because of the initially low uptake; less than ten percent of people in the control turned on targeting cookies. This means that the substantial improvement in the two relevant nagging conditions approximately doubled the number of people with that setting enabled.

resulted in greater uptake. Personalized ads χ^2 (1, N = 756) = 3.59, p = .06; Targeting cookies χ^2 (1, N = 756) = 16.73, p < .001.

[137] χ^2 (2, N = 306) = 1.47, p = .48.

[138] Note that the total percentile here is not the simple sum of the first and second nag percentiles as the denominator for the second nag will always be smaller than the one for the first nag so long as the first nag changes anyone's answer.

Again, however, there was no difference between the two relevant nagging conditions (recall that there was no nag here in the single-nag condition).[139]

Table 4: Nagging Type Effectiveness – Targeting Cookies

Accepted Targeting Cookies

Condition	Initially	After Nag	Total	Percent nagged who changed
Control	9.60%		9.60%	
Repeated Nags, No Interface Interference	8.63%	8.60%	17.23%	9.44%
Repeated Nags, Interface Interference	10.38%	13.66%	24.04%	15.24%
Single Nag, Interface Interference	10.05%		10.05%	

These results show that nags work even absent the other dark patterns. Further, they hint that nags have some diminishing returns and that the right nag, coming late, may be as potent as repeated earlier nags.

Due to the limited individual difference results in the first study, a more targeted approach was used here. We ran binary logistic regressions looking at the effects of the previously used tech skills and authoritarianism measures, along with a new privacy values measure on overall responses to the nags.[140] Specifically, these analyses looked only at people who received the nags and used whether those participants were successfully nagged as the dependent measure. For the subscription nags, this combined the results of nags 1 and 2 (in the conditions that had two nags). This analysis takes advantage of a unique feature of nagging dark patterns — each person is, in a sense, their own control. We know what each participant said before the dark pattern (namely "no") and we can then see whether, after the dark pattern, they now say "yes."

---

[139] χ^2 (1, N = 344) = 2.69, p = .10.

[140] The technology skills and authoritarianism scales were exactly the same as in our prior study, and each exhibited acceptable reliability (authoritarianism Cronbach's alpha = .87; tech skills Cronbach's alpha = .68). The privacy values questions were "I care a lot about whether the information I share with websites and apps remains private" and "when I create a new account on a website, I try to choose privacy protective settings." They also exhibited acceptable reliability (Cronbach's alpha = .68).

For personalized ads, nags were more likely to be successful for people who score high on the authoritarianism personality measure or low on the technology skills one. Privacy values did not have a significant effect.[141] The pattern was similar on susceptibility to the targeting cookies nag, with authoritarian personality increasing the nag's effectiveness, technology skills reducing it, and privacy values non-significantly reducing it as well.[142]

Though no dark patterns were used in the second half of the study, the matched identifiers, behavioral targeting, and Do Not Sell questions were still asked. As in the first study, a supermajority of those participants in this follow-up actively turned on "Do not sell or share my personal information."[143]

CONCLUSION

Though prior experimental work establishes that dark patterns are effective in prompting consumers to purchase goods and services they do not want, and that dark patterns can substantially shape the choices that consumers make when confronted with cookie consent screens/consent management platforms, there was an open question of whether dark patterns can also manipulate consumers into making privacy choices and adopting privacy settings that are contrary to their interests and preferences. Here we show that they can. Several dark patterns influenced participants as they completed an account set-up procedure that closely mirrored what consumers might encounter if signing up for a new video streaming service like Netflix, Hulu, or Peacock. This paper strongly suggests that dark patterns do prompt consumers to surrender more privacy than they otherwise would. We also show, for the first time, that nagging dark patterns, which manipulate consumers but do not deceive them, are highly effective even when used

---

141 Authoritarianism B = 0.54, Wald = 9.67, p < .001, Exp(B) = 1.72, 95% CI [1.22, 2.42].
Tech skills B = -0.50, Wald = 3.53, p = 0.06, Exp(B) = 0.61, 95% CI [0.36, 1.02].
Privacy values B = -0.22, Wald = 1.18, p = 0.28, Exp(B) = 0.8, 95% CI [0.54, 1.19].

142 Authoritarianism B = 0.51, Wald = 6.74, p = 0.01, Exp(B) = 1.66, 95% CI [1.13, 2.44].
Tech skills B = -0.59, Wald = 4.56, p = 0.03, Exp(B) = 0.56, 95% CI [0.32, 0.95].
Privacy values B = -0.35, Wald = 3.31, p = 0.07, Exp(B) = 0.7, 95% CI [0.48, 1.03].

143 Across all conditions, 71.0% chose this option. This did not differ significantly depending on prior nags (p = .46).

sparingly. These nagging dark patterns convince users to adopt privacy settings consumers do not prefer, not by persuading consumers to change their preferences, but by making it clear that the user interface will not take no for an answer. Of course, once the interface receives the answer its designers prefer, it will never prompt users to reconsider their choice.

The differences between the control group and the treatment groups subjected to various kinds of dark patterns are substantial and stark. Moreover, in our sample, dark patterns significantly and adversely affect all kinds of Americans — the rich and the poor, the old and the young, men and women, and both highly-educated and less-educated people. Among the different demographic characteristics we studied, only less technologically literate people and those exhibiting more authoritarian personality dispositions stood out as especially vulnerable, and those effects were somewhat inconsistent.

Notwithstanding the scholarly consensus from a variety of experimental papers suggesting that dark patterns can be quite effective at manipulating consumers into making choices that are inconsistent with their preferences, some academic and industry voices have argued against legislative and regulatory intervention on the grounds that consumers eventually will learn about dark patterns through repeated exposure and adopt effective self-defense mechanisms. With so many legislatures and regulators around the world considering the imposition of new limits on the technology sector, this debate is highly relevant in contemporary policy.

Our results show consumers cannot fully overcome dark patterns even when they try. While dark patterns are more effective at manipulating consumers who are making the choices they'd ordinarily make, they also thwart many consumers trying to choose the most privacy-protective options. Despite several years' worth of exposure to dark patterns, many consumers have not learned to defeat them. Indeed, it is plausible that consumers' exposure to dark patterns has created a kind of learned helplessness, where consumers conclude that they will eventually be manipulated into surrendering personal information they wish to keep private, so they figure they may as well surrender the information sooner rather than later, to save themselves from perpetually clicking "Maybe later" or "Stay signed out." Left to their own devices, consumers are frequently unable to navigate user interfaces that place dark patterns in the path of making privacy protective choices. Measures to

empower consumers to help themselves probably need to take the form of technological interventions, such as plug-ins or AI tools designed to counteract dark patterns in real time, rather than consumer learning or public awareness campaigns.[144] Our findings substantially strengthen the case for legislative or regulatory interventions to address the dark patterns problem. These interventions may take the form of prohibitions on dark pattern interfaces and/or mandates that web sites and apps respect browser-based privacy preference signals (such as Global Privacy Controls[145]), but it is becoming increasingly apparent that some sort of legal protections are likely necessary to enable consumers to exercise autonomous choice in online environments.

---

[144] *See* Jieshan Chen et al., *Unveiling the Tricks: Automated Detection of Dark Patterns in Mobile Applications*, 36 ANN. ACM SYMP. ON USER INTERFACE SOFTWARE AND TECH. (2023) (developing an automatic dark pattern detection system and testing its efficacy as an aid to mobile app users); Than Htut Soe et al, *Automated Detection of Dark Patterns in Cookie Banners: How to Do It Poorly and Why It is Hard to do It Any Other Way* (2022), https://arxiv.org/pdf/2204.11836 [https://perma.cc/EBG3-MLDQ] (discussing the challenges involved with using machine learning to detect dark patterns); Ioannis Stavrakakis et al., *A Framework of Web-Based Dark Patterns that can be Detected Manually or Automatically*, 14 J. ADVANCES INTERNET TECH. 36 (2001) (showing that while some dark patterns can be detected automatically, automatic detection of other kinds of dark patterns is quite difficult).

[145] *See* GLOBAL PRIVACY CONTROL, https://globalprivacycontrol.org/, (last visited Mar. 18, 2025) [https://perma.cc/FUY3-J7R2]. The CCPA treats websites that respect Global Privacy Control signals as having complied with CCPA's opt-out requirements. Thus, firms that honor Global Privacy Control requests do not need to create a "Do not sell or share my personal information" hyperlink on their landing pages. *See* CAL. CIV. CODE § 1798.185 (a)(19) (authorizing the CPPA to issue regulations on browser-based opt-out mechanisms); CAL. CODE REGS. TIT. 11 § 7025(g) (2025) (providing that browser-based opt-out preference signals are an alternative way to satisfy CCPA's opt-out requirements).

APPENDIX 1: SAMPLE DEMOGRAPHICS

	Study 1	Follow-up	Census[146]
Gender			
Female	51.5%	51.2%	50.8%
Male	48.5%	48.5%	49.2%
Other	.1%	.3%	
Age (Years)			
Median	41	40	
Mean	42.38 (17.68)	41.50 (13.74)	
Political Orientation (1–5)[147]	2.70 (1.18)	2.62 (1.10)	
Race and Ethnicity			
White alone	76.3%	81.3%	75.5%
Black or African American alone	12.6%	8.2%	13.6%
American Indian or Native American alone	0.7%	.4%	1.3%
Asian American alone	4.4%	4.5%	6.3%
Hawaiian or Pacific Islander alone	.1%	.1%	0.3%
Other or multiracial	5.9%	5.4%	3.0%
Hispanic (of any race)	14.6%	16.9%	19.1%
Educational Attainment			
Less Than High School Diploma	1.7%	0.8%	8.8%
High School Diploma or GED	26.4%	22.9%	28.5%
Two-Year or Some College	28.4%	32.0%	25.0%
Four-Year College	28.4%	30.2%	23.4%
Graduate Degree	15.1%	14.2%	14.2%

Note: For age and political orientation, standard deviation is in parentheses.

[146] Ethnicity and gender statistics are from the U.S. Census website. *QuickFacts*, U.S. CENSUS BUREAU, (last visited June 4, 2024) https://www.census.gov/quickfacts//fact//table//US//PST045217 [https://perma.cc/]; *Educational Attainment in the United States: 2022*, U.S. CENSUS BUREAU (Feb. 16, 2023), https://www.census.gov/data/tables/2022/demo/educational-attainment/cps-detailed-tables.html [https://perma.cc/UC9Z-S2D7] (educational attainment was calculated for the population over the age of 25 from data in table 1).

[147] Political orientation was assessed on a scale ranging from 1 (very liberal) to 5 (very conservative).

APPENDIX 2: ADDITIONAL DARK PATTERN SCREENS

2.1. Subscription nagging dark patterns in Study 1

Initial nag

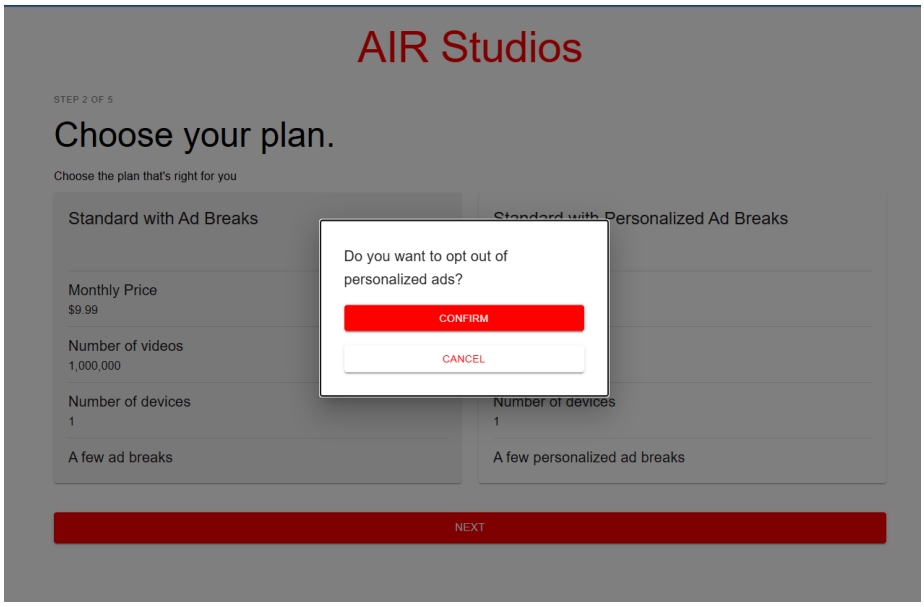

Second nag

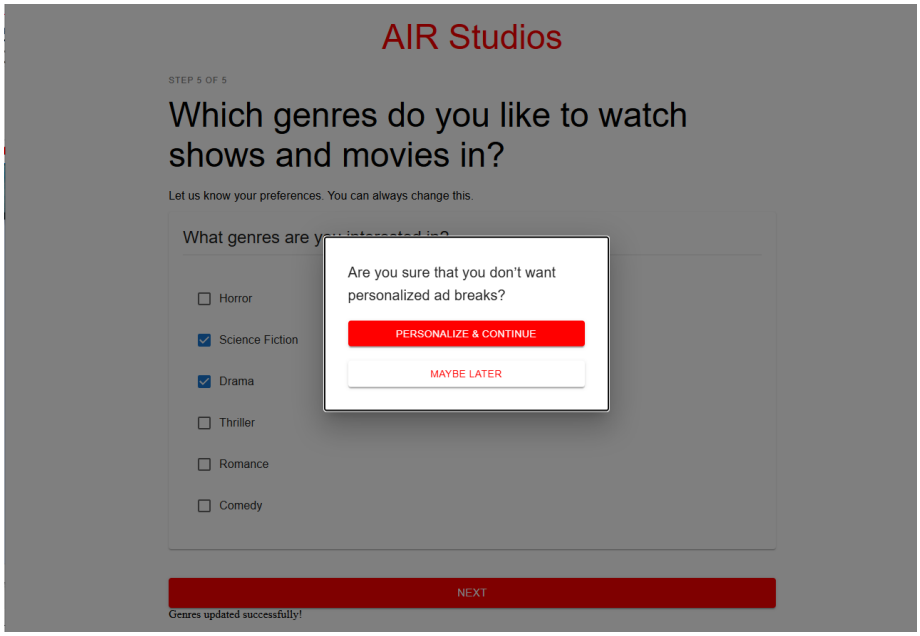

2.2. Targeting cookie nagging dark pattern in Study 1.

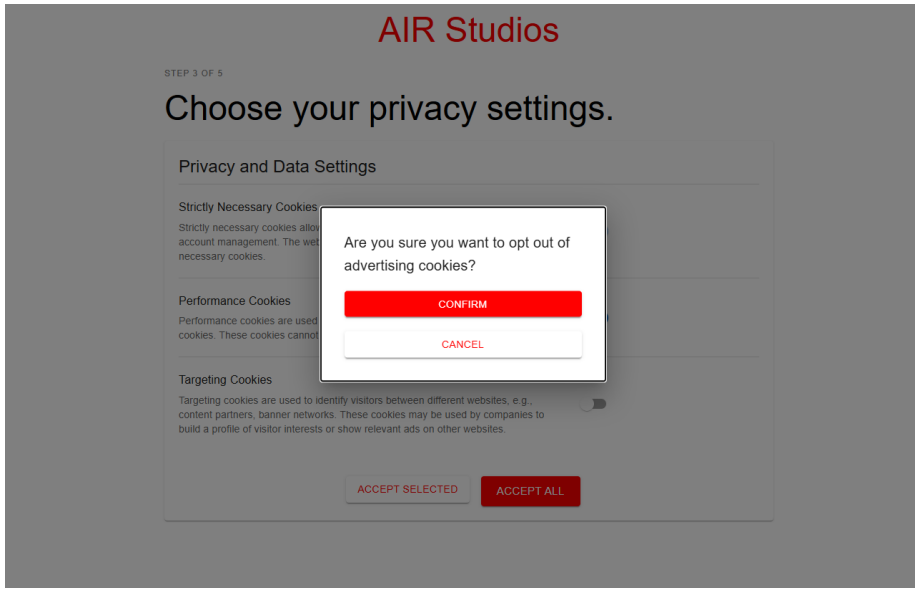

APPENDIX 3: NAGGING DARK PATTERN SCREENS

3.1. Subscription nagging dark patterns

Initial nag

No interface interference

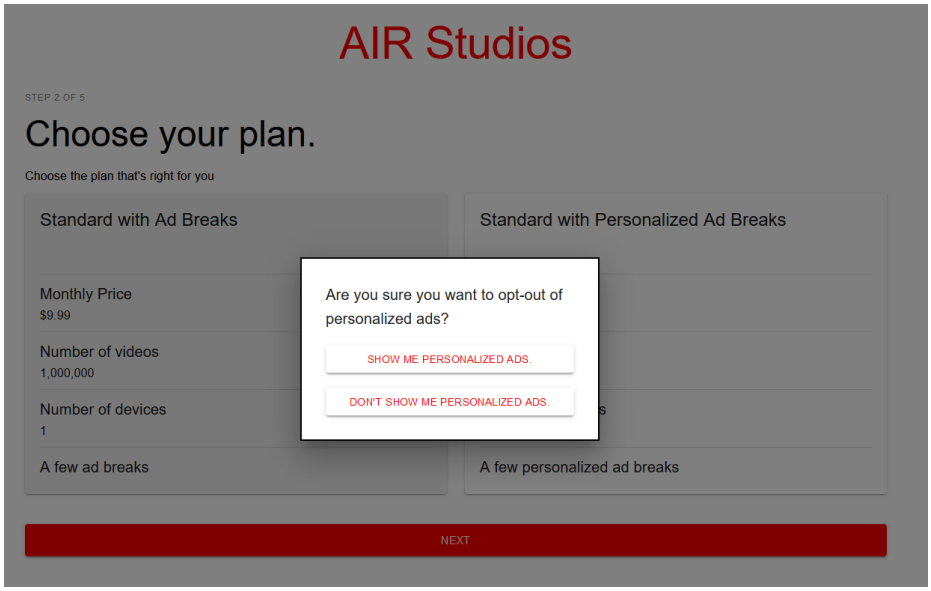

Question box with interface interference

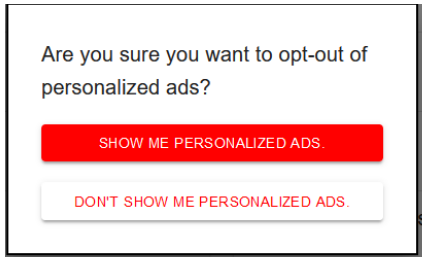

Second nag

No interface interference

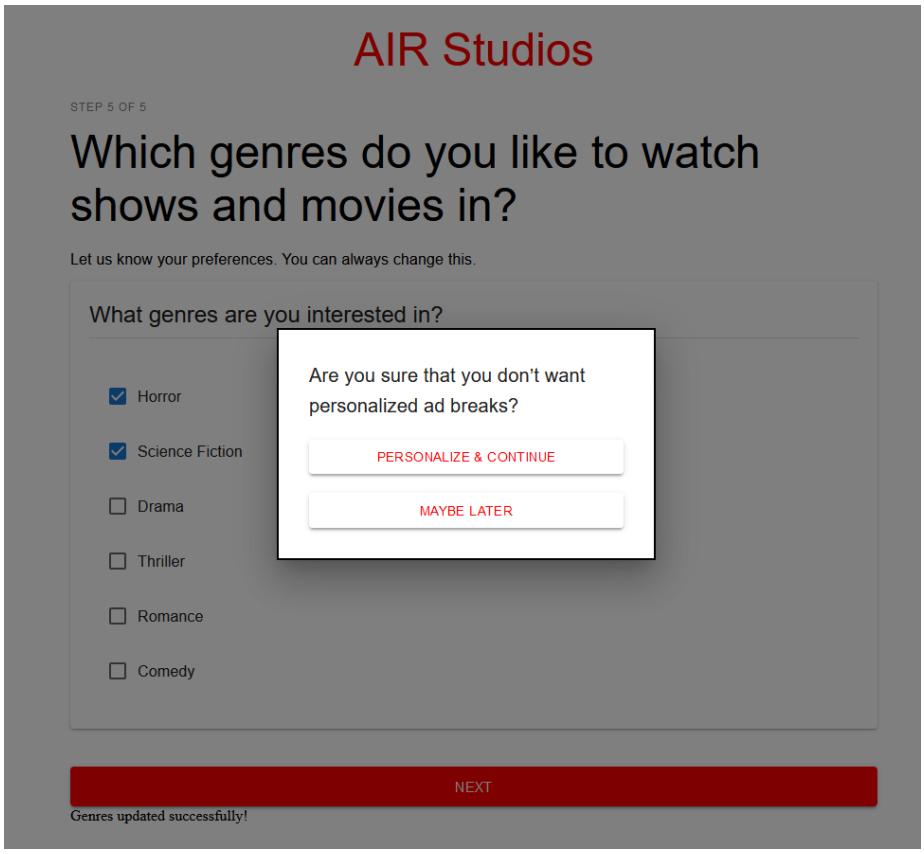

Question box with interface interference

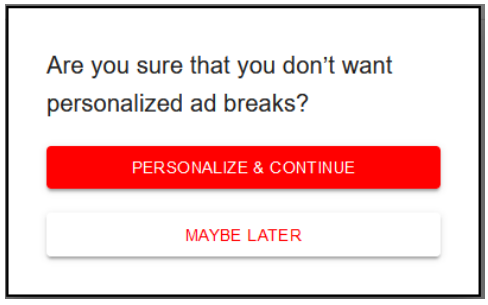

3.2. Targeting cookie nagging dark pattern.

No interface interference

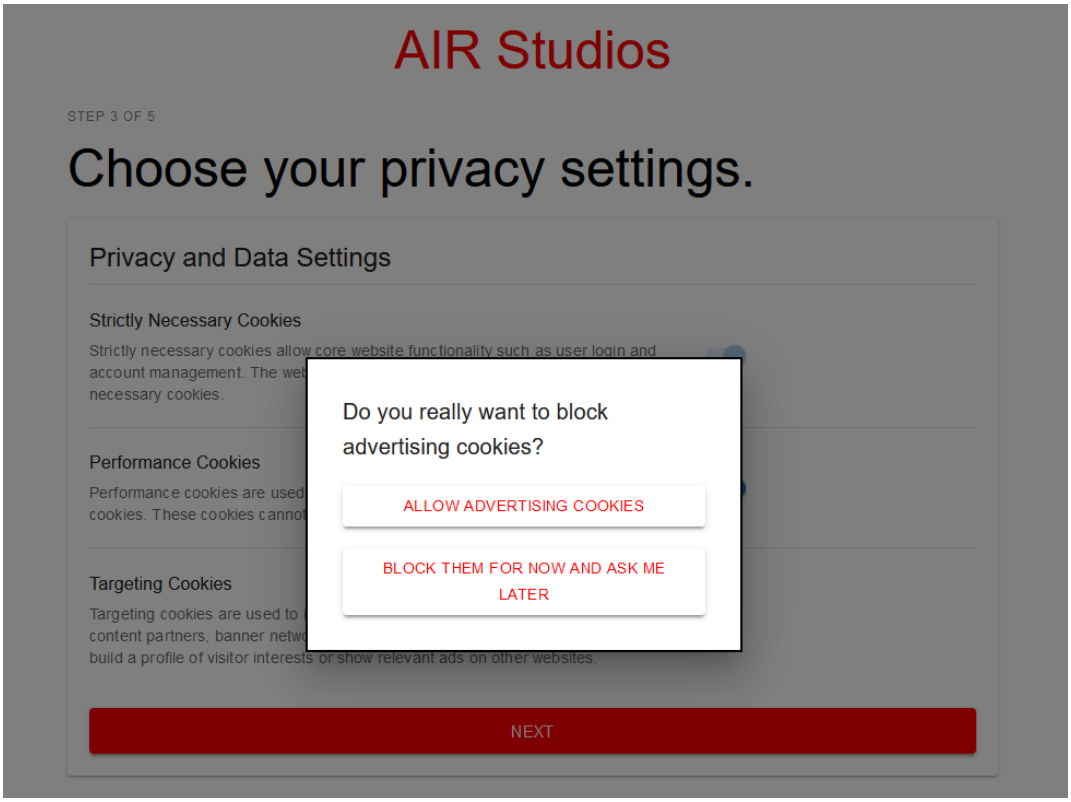

Question box with interface interference