

# Shining a Light on Dark Patterns

Jamie Luguri\* & Lior Jacob Strahilevitz\*\*

## Abstract

Dark patterns are user interfaces whose designers knowingly confuse users, make it difficult for users to express their actual preferences, or manipulate users into taking certain actions. They typically exploit cognitive biases and prompt online consumers to purchase goods and services that they do not want, or to reveal personal information they would prefer not to disclose. Research by computer scientists suggests that dark patterns have proliferated in recent years, but there is no scholarship that examines dark patterns' effectiveness in bending consumers to their designers' will. This article provides the first public evidence of the power of dark patterns. It discusses the results of the authors' large-scale experiment in which a representative sample of American consumers were randomly assigned to a control group, a group that was exposed to mild dark patterns, or a group that was exposed to aggressive dark patterns. All groups were told they had been automatically enrolled in an identity theft protection plan, and the experimental manipulation varied what acts were necessary for consumers to decline the plan. Users in the mild dark pattern condition were more than twice as likely to remain enrolled as those assigned to the control group, and users in the aggressive dark pattern condition were almost four times as likely to remain enrolled in the program. There were two other striking findings. First, whereas aggressive dark patterns generated a powerful backlash among consumers, mild dark patterns did not – suggesting that firms employing them generate substantial profits. Second, less educated subjects were significantly more susceptible to mild dark patterns than their well-educated counterparts. Both findings suggest that there is a particularly powerful case for legal interventions to curtail the use of mild dark patterns.

The article concludes by examining legal frameworks for ameliorating the dark patterns problem. Many dark patterns appear to violate federal and state laws restricting the use of unfair and deceptive practices in trade. Moreover, in those instances where consumers enter into contracts after being exposed to dark patterns, their consent could be deemed voidable under contract law principles. The article proposes a quantitative bright-line rule for identifying impermissible dark patterns. Dark patterns are presumably proliferating because firms' secret and proprietary A-B testing has revealed them to be profit maximizing. We show how similar A-B testing can be used to identify those dark patterns that are so manipulative that they ought to be deemed unlawful.

---

\* Law Clerk to the Honorable Brenda Sannes, United States District Court, Northern District of New York. J.D. University of Chicago Law School, 2019; Ph.D. Social Psychology, Yale University, 2015. The views expressed herein are solely those of the authors.

\*\* Sidley Austin Professor of Law, University of Chicago. For helpful comments on earlier drafts and conversations the authors thank Omri Ben-Shahar, Sebastian Benthall, Adam Chilton, Brett Frischmann, Meirav Furth-Matkin, Todd Henderson, William Hubbard, Filippo Lancieri, Anup Malani, Florencia Marotta-Wurgler, Jonathan Masur, Jonathan Mayer, Richard McAdams, Terrell McSweeney, Paul Ohm, Roseanna Sommers, Geof Stone, Kathy Strandburg, Cass Sunstein, Blase Ur, Mark Verstraete, and Luigi Zingales, along with workshop participants at the University of Chicago's PALS Lab, the University of Chicago Works in Progress Workshop, and the Stigler Center's 2019 Antitrust and Competition Conference. The authors thank the Carl S. Lloyd Faculty Fund for research support and Tyler Downing and Daniel Jellins for excellent research assistance.

## Table of Contents

Introduction.....	3
I. Dark Patterns in the Wild .....	6
II. An Experimental Test of the Effectiveness of Dark Patterns .....	17
A. Rates of Acceptance.....	21
B. The Influence of Stakes .....	23
C. Potential Repercussions of Deploying Dark Patterns .....	24
D. Predicting Dark Pattern Susceptibility .....	27
III. Are Dark Patterns Unlawful?.....	29
A. Laws Governing Deceptive and Unfair Practices in Trade .....	30
B. Other Relevant Federal Frameworks .....	37
C. Contracts and Consent .....	38
D. Line Drawing .....	43
E. Persuasion .....	46
Conclusion .....	48

## Introduction

Everybody has seen them before and found them frustrating, but most consumers don't know what to call them. They are what computer scientists have (for the last decade) described as *dark patterns*,<sup>1</sup> and they are a proliferating species of sludge (to use a term preferred by behavioral economists)<sup>2</sup> or market manipulation (the moniker preferred by some legal scholars).<sup>3</sup> Dark patterns are user interfaces whose designers knowingly confuse users, make it difficult for users to express their actual preferences, or manipulate users into taking certain actions. They typically prompt users to rely on System 1 decision-making rather than more deliberate System 2 processes, exploiting cognitive biases like framing effects, the sunk cost fallacy, and anchoring. The goal of most dark patterns is to manipulate the consumer into doing something that is inconsistent with her preferences, in contrast to marketing efforts that are designed to alter those preferences. The first wave of academic research into dark patterns identified the phenomenon and developed a typology of dark pattern techniques.<sup>4</sup>

This summer, computer scientists at Princeton and the University of Chicago took a second step towards tackling the problem by releasing the first major academic study of the prevalence of dark patterns.<sup>5</sup> Arunesh Mathur and six co-authors developed a semi-automated method for crawling more than 11,000 popular shopping websites. Their analysis revealed the presence of dark patterns on more than 11% of those sites, and the most popular sites were also most likely to employ dark patterns.<sup>6</sup>

If the first wave of scholarship created a useful taxonomy and the second step in the scholarship established the growing prevalence of dark pattern techniques then it seems clear where the literature ought to go next. Scholars need to quantify the effectiveness of dark patterns in convincing online consumers to do things that they would otherwise prefer not to

---

<sup>1</sup> User interface designer Harry Brignull coined the phrase in 2009 and maintains a web site that documents them in an effort to shame the programmers behind them. See <https://www.darkpatterns.org/>

<sup>2</sup> Cass R. Sunstein, *Sludges and Ordeals*, 69 DUKE L.J. \_\_\_\_ (forthcoming 2019); Richard H. Thaler, *Nudge, Not Sludge*, 361 SCIENCE 431 (2018). A sludge is an evil "nudge," one that exploits their cognitive biases to persuade them to do something that is undesirable, typically by introducing excessive friction into choice architecture. See Cass R. Sunstein, *Sludge Audits*, Harvard Public Law Working Paper No. 19-21 (July 2, 2019 draft), available at [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3379367](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3379367) (hereinafter Sunstein, *Sludge Audits*).

<sup>3</sup> Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995 (2014); Jon D. Hanson & Douglas A. Kysar, *Taking Behavioralism Seriously: The Problem of Market Manipulation*, 74 NYU L. REV. 632 (1999).

<sup>4</sup> See, e.g., Christoph Bösch et al., *Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns*, *Proceedings on Privacy Enhancing Technologies* (2016); Colin M Gray et al., *The Dark (Patterns) Side of UX Design*, *PROCEEDINGS OF THE 2018 CHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS* (2018).

<sup>5</sup> Arunesh Mathur et al., *Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites*, July 17, 2019 working paper, available at <https://webtransparency.cs.princeton.edu/dark-patterns/assets/dark-patterns-v2.pdf>.

<sup>6</sup> As the authors themselves note, see *id.* at 2, this figure probably understates the prevalence of dark patterns, because their taxonomy of dark patterns leaves out several important dark pattern mechanisms, perhaps in part because they are hard to identify using the semi-automated approach employed by the authors. See *infra* table 1 (providing a taxonomy that omits nagging, bait and switch, aesthetic manipulation and other important types of dark patterns).

do. In short, the question we pose in this paper is “how effective are dark patterns?” That is not a question that has been answered in academic research to date. But it is a vital inquiry if we are to understand the magnitude of the problem and whether regulation is appropriate.

To be sure, the lack of published research does not mean that the effectiveness of these techniques is a complete mystery. On the contrary, we suspect that the kind of research results we report here have been replicated by social scientists working in-house for technology and ecommerce companies. Our hunch is that consumers are seeing so many dark patterns in the wild because the internal, proprietary research suggests dark patterns are presently profit-maximizing for the firms that employ them. But those social scientists have had strong incentives to suppress the results of their A-B testing of dark patterns, so as to preserve data about the successes and failures of the techniques as trade secrets and (perhaps) to stem the emergence of public outrage and legislative or regulatory responses.

With bipartisan legislation that would constrain the use of dark patterns currently pending in the Senate,<sup>7</sup> investigative reporters beginning to examine the dark patterns problem,<sup>8</sup> and one of the nation’s leading privacy scholars testifying before the Federal Trade Commission (F.T.C.) that in his estimation, 2019 would be the year of the dark pattern,<sup>9</sup> e-commerce firms probably expect that, where the effectiveness of dark patterns is concerned, heat will follow light. So they have elected to keep the world in the dark for as long as possible. The strategy has worked so far.

The basic problem of manipulation in marketing and sales is not unique to interactions between computers and machines. The main factors that make this context interesting are its relative newness and scale. In both traditional and online contexts legal actors have to make tough decisions about where the precise line is between persuasion and manipulation, and what conduct is misleading enough to eliminate what might otherwise be constitutionally protected rights for sellers to engage in commercial speech. The law has long elected to prohibit certain strategies for convincing people to part with money or personal information. Laws prohibiting fraud have been around, seemingly forever, and more recently implemented laws proscribe pretexting. States and the federal government have given consumers special rights in settings characterized by high pressure, mild coercion, or vulnerability, such as door-to-door-sales, and transactions involving funeral services, timeshares, telemarketing, or home equity loans. Sometimes the law enacts outright prohibitions with substantial penalties. Other times it creates cooling off periods that cannot be waived. A key question we address is what online tactics are egregious enough to warrant this kind of special skepticism.

---

<sup>7</sup> Deceptive Experiences to Online Users Reduction Act (DETOUR Act), Senate Bill 1084, 116th Congress, introduced April 9, 2019, text available at <https://www.congress.gov/bill/116th-congress/senate-bill/1084/text>.

<sup>8</sup> See, e.g., Jennifer Valentino-Devries, *How E-Commerce Sites Manipulate You into Buying Things You May Not Want*, N.Y. TIMES, June 24, 2019, at B1.

<sup>9</sup> See F.T.C. Hearing, Competition and Consumer Protection in the Twenty-First Century, April 9, 2019, Testimony of Professor Paul Ohm, Georgetown University, Transcript at 49 (“my prediction for 2019 ... is this is the year where dark patterns really becomes the kind of thing that we’re really talking a lot about.”), available at [https://www.FTC.gov/system/files/documents/public\\_events/1418273/FTC\\_hearings\\_session\\_12\\_transcript\\_day\\_1\\_4-9-19.pdf](https://www.FTC.gov/system/files/documents/public_events/1418273/FTC_hearings_session_12_transcript_day_1_4-9-19.pdf).

Ours is a descriptive paper, an empirical paper, a normative paper, and then a descriptive paper again. That said, the new experimental data we reveal is the most important take-away. Part I begins by describing dark patterns – what techniques they include and what some of the most prominent examples are. The description illuminates several real-world dark patterns and suites of dark patterns employed by major multinational corporations. The Part also provides a streamlined taxonomy of dark pattern techniques, one that builds on work that computer scientists have done while providing some conceptual clarity that’s caused scholars of human-computer interactions to lump together divergent phenomena.

Part II provides the paper’s core contribution. As scholars have seen the proliferation of dark patterns, many have assumed that dark patterns are efficacious. Why else would large, well-capitalized companies that are known to engage in A-B testing be rolling them out? Judges confronting dark patterns have for the most part shared these intuitions, though not universally. We show that many widely employed dark patterns prompt consumers to do what they would not do in a more neutral decision-making environment. But beyond that, we provide the first comparative evidence that quantifies how well they work, and that sheds some light on the question of which techniques work best. Our bottom line is that *dark patterns are strikingly effective in getting consumers to do what they would not do when confronted with more neutral user interfaces*. Relatively mild dark patterns more than doubled the percentage of consumers who signed up for a dubious identity theft protection service that we told our subjects we were selling, and aggressive dark pattern nearly quadrupled the percentage of consumers signing up. In social science terms, the magnitudes of these effects are enormous. We then provide powerful evidence that dosage matters – aggressive dark patterns generate a powerful customer backlash. Mild dark patterns usually do not, and therefore, counterintuitively, the strongest case for regulation and other legal intervention concerns subtle uses of dark patterns. Finally, we provide compelling evidence that less educated Americans are significantly more vulnerable to dark patterns than their more educated counterparts, and that trend is particularly pronounced where subtler dark patterns are concerned. This observation raises distributive issues and is also useful as we consider how the law might respond to dark patterns.

Part III looks at the existing law and asks whether it prohibits dark patterns. This is an important area for inquiry because pending bipartisan legislation proposes that the F.T.C. be given new authority to prohibit dark patterns.<sup>10</sup> It turns out that with respect to a number of central dark pattern techniques, the F.T.C. is already going after some kinds of dark patterns, and the federal courts have been happy to cheer the agency along. The most successful actions have nearly all fallen under the F.T.C.’s section five authority to regulate deceptive acts and practices in trade. To be sure, other important dark patterns fit less comfortably within the categories of deceptive or misleading trade practices, and there is lingering uncertainty as to how much the F.T.C.’s authority to restrict unfair trade practices will empower the agency to restrict that behavior. The passage of federal legislation aimed squarely at dark patterns would provide useful new legal tools, but there is no reason to delay enforcement efforts directed at egregious dark patterns while waiting on Congress to do something.

Of course, the F.T.C. lacks the resources to be everywhere, so a critical issue going forward will be whether contracts that are agreed to in large measure because of a seller’s use of dark patterns are deemed valid. This issue is just now starting to bubble up in the case law. To deal with this question, and other important line-drawing questions, we propose a quantitative

---

<sup>10</sup> See *supra* text accompanying note 7.

“more likely than not” approach to regulation. The method we use in this paper is easy to replicate, and the math is not especially fancy. Where the use of a dark pattern technique more than doubles the rate of acceptance compared to neutral choice architecture, the law should regard the dark pattern’s use as per se unlawful. To be sure, that is an underinclusive test, one that should be supplemented by a balancing inquiry. But we think it is a good and straightforward place to start as the law begins to grapple seriously with the question of how to regulate dark patterns. Notably, both the mild and aggressive dark patterns we tested experimentally satisfied that test. As we explain in Part III, there is a plausible case to be made that agreements procured through the use of dark patterns are voidable as a matter of contract law under the undue influence doctrine.

We said at the outset that dark patterns are different than other forms of dodgy business practices because of the scale of e-commerce. There may be poetic justice in the notion that this very scale presents an opportunity for creative legal regulators. It is exceedingly difficult to figure out whether a door to door salesperson’s least savory tactics significantly affected the chances of a purchase – was the verbal sleight of hand material or incidental? Who knows? But with e-commerce, firms can run thousands of consumers through identical interfaces at a reasonable cost and see how small tweaks to the software might alter user behavior. Social scientists working in academia or for the government can do this too; we just haven’t done so before today. Now that scholars can test dark patterns, we can isolate causation in a way that’s heretofore been impossible in the brick-and-mortar world. Unlike brick-and-mortar manipulation, dark patterns are hiding in plain sight, operate on a massive scale, and are relatively easy to detect. Those facts strengthen the case further for the legal system to address their proliferation.

So let’s spend some time getting to know dark patterns.

## I. Dark Patterns in the Wild

Suppose you are getting commercial emails from a company and wish to unsubscribe. If the company is following the law they will include in their emails a link to a page that allows you to remove your email address.<sup>11</sup> Some companies make that process simple, automatically removing your address when you click on an unsubscribe link or taking you to a page that asks you to type in your email address to unsubscribe. Once you do so they will stop sending you emails.

Other companies will employ various tools to try to keep you on their lists. They may remind you that if you unsubscribe you will lose out on valuable opportunities to save money on their latest products (dark patterns researchers call this practice “confirmshaming”). Or they’ll give you a number of options besides the full unsubscribe that most people presumably want, such as “receive emails from us once a week” or “receive fewer emails from us” while making users who want to receive no more emails click through to a subsequent page.<sup>12</sup> (These

---

<sup>11</sup> This is required by the CAN-SPAM Act of 2003, 15 U.S.C. § 103.

<sup>12</sup> As of July 2019, Best Buy’s unsubscribe link in commercial emails followed this pattern. If a user clicked on the unsubscribe hyperlink at the bottom of a marketing email, she would be taken to a screen that provided three options: “Receive all General Marketing emails from Best Buy.” [This box is checked by default, so a user who clicks “unsubscribe” and then “submit” will not stop receiving emails from Best Buy.] The second option says, “Receive no more than one General Marketing email per week.” And the third option is “Receive no General Marketing emails (unsubscribe).”

techniques are referred to as “obstruction” dark patterns).<sup>13</sup> The company is making it easy for you to do what it prefers (you continue to receive lots of marketing emails), and harder for you to do the thing it can live with (receiving fewer emails), or the thing you probably prefer and are entitled to by law (receiving no emails from the company).

In other instances, firms employ highly confusing “trick question” prompts that make it hard for even smart consumers to figure out how they are to accomplish their desired objective. For instance, see the membership cancellation page from the Pressed Juicery:<sup>14</sup>

## Membership Status

---

### Canceling your membership?

Are you sure you want to cancel your membership? You will no longer receive membership pricing on all our products.

**CONTINUE**

**CANCEL**

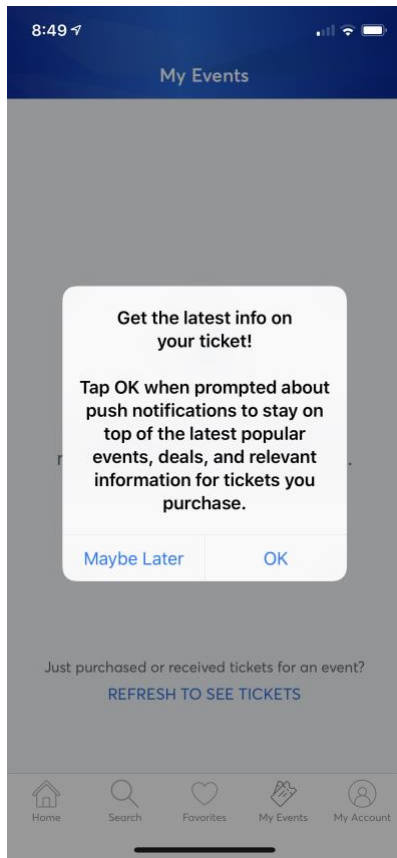
Other aggravating examples of dark patterns abound. If you have found it easy to sign up for a service online, with just a click or two, but when it came time to cancel the service had to make a phone call or send a letter via snail mail, you have been caught in a “roach motel” dark pattern (it’s easy to get in but hard to get out). If you’ve ever seen an item in your shopping cart that you did not add to it and wondered how it got there, you have encountered a “sneak into the cart” dark pattern. If you’ve once been given a choice between signing up for notifications, with the only options presented being “Yes” and “Not Now,” only to be asked again about signing up for notifications two weeks later when you select “Not Now,” that’s a “nagging” dark pattern. Here is one from Ticketmaster’s smartphone app.<sup>15</sup>

---


<sup>13</sup> Gray et al., *supra* note 4, at 5-6; LIOR STRAHILEVITZ ET AL., SUBCOMMITTEE REPORT: PRIVACY AND DATA PROTECTION, STIGLER CENTER COMMITTEE FOR THE STUDY OF DIGITAL PLATFORMS 22-23 (2019).

<sup>14</sup> For this example, we thank Karin Curkowicz. See <https://twitter.com/KCurkowicz/status/1137855721507213312>

<sup>15</sup> Google Maps does essentially the same thing. When a user repeatedly travels to a particular location and uses the app’s directions, the app will display a “Go here often?” pop-up window that asks whether the location is her “Home,” “Work,” or “Other” (school, gym, etc.) approximately once a week. A user can close the window each time but there is evidently no way to prevent the queries from reappearing short of deleting all location history. The pop-up window notes that users’ labels for locations “will be used across Google products, for personalized recommendations, and for more useful ads.”



Bait and switch is another time-tested dodgy business practice, and the tactic has emerged online as a type of dark pattern. Sometimes it arises in its classic form, and sometimes it emerges as a bait-and-sell and switch, where the customer does get to purchase the good or service that was advertised, but is then shown a barrage of ads for things the customer does not want. Here is an example of the latter from one of the author's recent online purchases from the aforementioned Ticketmaster.





Alexander Hamilton was generally depicted clean-shaven in portraits.<sup>16</sup> Other than that, it's not clear what the connection is between the ticket purchase and a razor-blade subscription. Notice further that besides bait and switch there are two subtler dark patterns embedded in the image above. In the ad, "Yes please" appears against a bright blue background while "No thanks" appears less legible against a gray one. Moreover, another even lighter gray font (barely legible in the pop-up ad) reveals important text about what a consumer has to click on to skip the several bait-and-switch ads that would follow, which in this case included further offers from Hotels.com, Priceline, and Hulu. The font appears much less prominently than the darker text above it about the razor blade offer.

Another common dark pattern is generating false or misleading messages about demand for products or testimonials. Arunesh Mathur and co-authors recently revealed that a number of popular shopping sites display information about recent sales activities that is driven by random number generators and similar techniques. For example, they caught thredup.com using a random number generator to display information about how many of a particular product were "just sold" in the last hour, and they found various sports jersey sales sites using identically phrased customer testimonials but with different customer names each time.<sup>17</sup> When a site notes that Anna in Anchorage just purchased a jacket that a user is examining, the academic research suggests these high-demand messages should be taken with a large grain of salt.

Having introduced a few vivid examples of dark patterns, it seems appropriate to introduce a workable taxonomy of the techniques. Several have been developed in the existing literature. One problem is that as interest in dark patterns has grown, so has the ostensible list of what counts as one. Putting together the various taxonomies in the literature results in a rather lengthy list, with some techniques being very problematic and others less so. There have been four key taxonomies to emerge in the dark patterns literature, with each building on and tweaking what came before. The chart below reproduces the current aggregated taxonomy in the literature and identifies which types of dark patterns have been identified in multiple taxonomies versus only some.<sup>18</sup> Our literature review reveals eight categories of dark patterns and 27 variants.<sup>19</sup> After presenting this information we will propose a modified, streamlined taxonomy that appropriately focuses on the means (the manipulative techniques used) rather than the ends (getting users to provide sensitive information, cash, recruit others, etc.). It is worth noting at the outset that some of the choices different teams of scholars have made in presenting their taxonomies relate to their different objectives. For example, some scholars, like Bösch et al, are not trying to be comprehensive. Others, like Mathur et al., are focusing on the sorts of dark patterns that can be identified using a semi-automated web-crawling process. Such processes lend themselves to flagging certain kinds of dark patterns (such as low-stock messages) more readily than others (such as toying with emotions).

---

<sup>16</sup> Alas, so was Aaron Burr.

<sup>17</sup> Mathur et al., *supra* note 5, at 18-19.

<sup>18</sup> Most of the dark patterns literature is co-authored. For space-saving reasons, we include in the table only the surname of the first listed author of such work.

<sup>19</sup> We apologize for the small font size necessary to squeeze the table onto a page. We promise the table is not intended to be a dark pattern – we actually want you to read the categories and examples closely.

*Table 1: Summary of Existing Dark Pattern Taxonomies*

Category	Variant	Description	Source
<b>Nagging</b>		Repeated requests to do something firm prefers	Gray
<b>Social Proof</b>	Activity messages	False / misleading Notice that others are purchasing, contributing	Mathur
	Testimonials	False / misleading positive statements from customers	Mathur
<b>Obstruction</b>	Roach Motel	Asymmetry between signing up and canceling	Gray, Mathur
	Price Comparison Prevention	Frustrates comparison shopping	Brignull, Gray, Mathur
	Intermediate Currency	Purchases in virtual currency to obscure cost	Brignull
	Immortal Accounts	Account and consumer info cannot be deleted	Bösch
<b>Sneaking</b>	Sneak into Basket	Item consumer did not add is in cart	Brignull, Gray, Mathur
	Hidden Costs	Costs obscured / disclosed late in transaction	Brignull, Gray, Mathur
	Hidden subscription / forced continuity	Unanticipated / undesired automatic renewal	Brignull, Gray, Mathur
	Bait & Switch	Customer sold something other than what's originally advertised	Gray
<b>Interface Interference</b>	Hidden information / aesthetic manipulation	Important information visually obscured	Gray
	Preselection	Firm-friendly default is preselected	Bösch, Gray
	Toying with emotion	Emotionally manipulative framing	Gray
	False hierarchy / pressured selling	Manipulation to select more expensive version	Gray, Mathur
	Trick questions	Intentional or obvious ambiguity	Gray, Mathur
	Disguised Ad	Consumer induced to click on something that isn't apparent ad	Brignull, Gray
	Confirmshaming	Choice framed in way that makes it seem dishonorable, stupid	Brignull, Mathur
	Cuteness	Consumers likely to trust attractive robot	Lacey
<b>Forced Action</b>	Friend spam / social pyramid / address book leeching	Manipulative extraction of information about other users	Brignull, Bösch, Gray
	Privacy Zuckering	Consumers tricked into sharing personal info	Brignull, Bösch, Gray
	Gamification	Features earned through repeated use	Gray
	Forced Registration	Consumer tricked into thinking registration necessary	Bösch
<b>Scarcity</b>	Low stock message	Consumer informed of limited quantities	Mathur
	High demand message	Consumer informed others are buying remaining stock	Mathur
<b>Urgency</b>	Countdown timer	Opportunity ends soon with blatant visual cue	Mathur
	Limited time message	Opportunity ends soon	Mathur

Now let's see if we can do a little bit of streamlining. In our view, dark patterns are techniques used to manipulate users to do things they would not otherwise do. Precisely what users wind up doing is irrelevant for our purposes, so long as it is something they do not genuinely want to do. This warrants removal of dark pattern techniques included above that are focused on ends rather than means.

Immortal accounts are a privacy-focused ostensible dark pattern, one that obstructs the deletion of information the customer may want to make disappear. Because the technique focuses on ends (privacy protection) rather than the mechanism used, we don't include it as a dark pattern. The same can be said about friend spam and privacy zuckering. Making robots cute to get people to share intimate details about themselves (which Lacey and Caudwell have dubbed a dark pattern)<sup>20</sup> is not appropriately characterized in that way. Part of the reason why is the ends-orientation identified above.<sup>21</sup> Gamification and non-misleading forms of forced registration are not dark patterns for different reasons. In our view, if a company wants to structure the quid pro quo that's central to their business model as "you give us personal information in exchange for stuff," this is permissible. So an online newspaper can decide to provide content for free in exchange for the user accurately identifying himself to facilitate subsequent marketing. As long as the nature of the exchange isn't concealed, it's not a dark pattern. So too with a business model that privileges highly engaged users over occasional ones. Finally, it seems to us that in the Mathur et al. framework "Scarcity" and "Urgency" are exploiting the same behavioral mechanisms to induce a type 1 purchase or disclosure decision. They can be collapsed for analytical purposes into a single category. Our edits produce the following revised taxonomy that is a bit easier on the eyes (and perhaps the brain).

---

<sup>20</sup> Cherie Lacey & Catherine Caudwell, *Cuteness as a 'Dark Pattern' in Home Robots*, in 14TH ACM/IEEE INT'L CONF. ON HUMAN-ROBOT INTERACTION (HRI) CONFERENCE PROCEEDINGS 374 (2019), available at <https://ieeexplore.ieee.org/abstract/document/8673274>.

<sup>21</sup> A deeper concern related to stretching the dark pattern label that far is the problem of identifying what a neutral baseline looks like. As we will argue below, if a neutral interface can be identified and then compared to an ostensibly manipulative one, we can use quantitative techniques to resolve lingering uncertainty about when a sales technique crosses the line. But it's hard to say what a "neutrally cute" robot looks like. More broadly, it has long been true that sellers of goods use conventionally attractive people to sell not only obvious products like fashion and jewelry but also less obvious products like detergent and insurance. The "cute robot" strategy is a variant of that. While it is possible to use A-B testing to identify the precise impact that a conventionally attractive model versus an average-looking model has on the sales of toothpaste, the lack of deception used in such advertisements and the extremely long pedigree of such techniques in advertising make this a poor fit for the category. Further, our prior is that the effect sizes from those techniques would be nontrivial but not especially large.

Table 2: Revised Taxonomy of Dark Patterns

Category	Variant	Description	Source
<b>Nagging</b>		Repeated requests to do something firm prefers	Gray
<b>Social Proof</b>	Activity messages	Misleading notice about other consumers' actions	Mathur
	Testimonials	Misleading statements from customers	Mathur
<b>Obstruction</b>	Roach Motel	Asymmetry between signing up and canceling	Gray, Mathur
	Price Comparison Prevention	Frustrates comparison shopping	Brignull, Gray, Mathur
	Intermediate Currency	Purchases in virtual currency to obscure cost	Brignull
<b>Sneaking</b>	Sneak into Basket	Item consumer did not add is in cart	Brignull, Gray, Mathur
	Hidden Costs	Costs obscured / disclosed late in transaction	Brignull, Gray, Mathur
	Hidden subscription / forced continuity	Unanticipated / undesired automatic renewal	Brignull, Gray, Mathur
	Bait & Switch	Customer sold something other than what's originally advertised	Gray
<b>Interface Interference</b>	Hidden information / aesthetic manipulation / false hierarchy	Important information visually obscured	Gray, Mathur
	Preselection	Firm-friendly default is preselected	Bösch, Gray
	Toying with emotion	Emotionally manipulative framing	Gray
	Trick questions	Intentional or obvious ambiguity	Gray, Mathur
	Disguised Ad	Consumer induced to click on something that isn't apparent ad	Brignull, Gray
	Confirmshaming	Choice framed in way that seems dishonest / stupid	Brignull, Mathur
<b>Forced Action</b>	Forced Registration	Consumer tricked into thinking registration necessary	Bösch
<b>Urgency</b>	Low stock / high-demand message	Consumer falsely informed of limited quantities	Mathur
	Countdown timer / Limited time message	Opportunity ends soon with blatant false visual cue	Mathur

That revised taxonomy of dark patterns is still lengthy, but it's hopefully easier to internalize. As we will see, in many instances firms are going to combine several of the techniques on this list. We previously singled out Ticketmaster, a major American company with a large market share. So we'll end our introduction to dark patterns with an extended exploration of techniques employed by Sony, a major Japanese company with a large market share.

As gaming platforms have become a major source of revenue, the dominant platforms have sought to profit off the increased appeal of online gaming. Online gaming allows individuals to play over the Internet against friends or strangers, and both Sony and Microsoft have made major investments in this technology. One strategy that is widely employed in popular games makes it necessary for players to sign up for the online platforms in order to earn the most appealing available rewards. One of the authors has a child who enjoys EA's FIFA soccer games on the Sony PlayStation, and to that end, the author signed up for a short-term subscription to PlayStation Plus – Sony's online gaming platform. During the next few paragraphs we will use Sony's user interface as a case study of dark patterns.

Let's begin with Sony's pricing model and graphic design choices.



Several notable aspects of the user interface stand out. First, notice the visual prominence of the 12 month subscription rather than the alternatives in the default view. This “false hierarchy” graphic design approach is a kind of dark pattern; one with a long and infamous history, at that.



Translation: "Do you agree with the reunification of Austria with the German Reich that was enacted on 13 March 1938 and do you vote for the party of our leader; Adolf Hitler?; Yes (large circle); No (small circle)"

Choice-architects have long understood that contrasting visual prominence can be used to nudge choosers effectively into a choice the architect prefers, and false hierarchy can come in handy whether one is an innovative multinational technology company or a group of genocidal fanatics conducting a sham election.<sup>22</sup> The visual contrast is one of the least subtle and presumably more benign dark patterns that can be encountered in the wild. Unlike many dark patterns identified above, it is almost certainly a legal marketing tactic when used in isolation.<sup>23</sup>

---

<sup>22</sup> It hopefully goes without saying that our juxtaposition is not equating the Sony Corporation with Nazi Germany.

<sup>23</sup> Most of the brick-and-mortar equivalent of dark pattern techniques on our list are either very uncommon or are widely believed to be dubious or unlawful when practiced in brick and mortar establishments. For example, telemarketers are prohibited from engaging in various forms of nagging, such as continuing to call someone who has said she does not wish to receive such calls. 10 C.F.R. 310.4(b)(1)(iii)(A). To take another example, the F.T.C. has issued a policy statement making it clear that the use of disguised ads is actionable under section 5. See F.T.C., Enforcement Policy Statement on Deceptively Formatted Advertisements (Dec. 22, 2015), available at [https://www.FTC.gov/system/files/documents/public\\_statements/896923/151222deceptiveenforcement.pdf](https://www.FTC.gov/system/files/documents/public_statements/896923/151222deceptiveenforcement.pdf). And some brick and mortar equivalents of dark patterns are so obviously unlawful that reputable firms do not even try them. For example, suppose Trader Joe's instructed its cashiers to start charging their customers for granola bars they did not purchase and slipping those bars into their shopping carts surreptitiously. Surely some customers who didn't notice what happened in the check-out aisle will not wind up returning the granola bars because of the inconvenience involved, but that hardly means there is no injury from the conduct, nor would we be comfortable describing the transaction as one in which the customers consented to purchase the granola bars. The online equivalent of that conduct is "sneak into basket." It's hard to imagine what a coherent defense of the tactic at a grocery store would look like.

Now let's consider Sony's pricing strategy. It is no mystery what Sony is trying to do. They want to minimize customer churn. Acquiring subscribers is quite costly, so Sony wants to retain them once they obtain them. Moreover, some customers may subscribe for a year because of the lower per-month rate (\$5 per month versus \$10 per month), and then grow bored with the service – these customers are good for Sony because they will be paying for network resources that they do not use, which will improve the experience for other paying customers. It's akin to people joining a gym as a New Year's resolution and then never showing up after January to use the treadmills. One way to get customers to commit to a longer term is by substantially discounting the monthly cost for customers who are willing to sign up for a year's membership. In this instance, Sony is willing to charge such customers half as much as customers who will only commit to subscribing for a month. To be clear, there is nothing legally wrong with Sony pursuing this pricing model and (at least from our perspective) there is not anything morally dubious about the practice either, at least not yet. The pricing model is not a dark pattern.

It's on the following screen that things get dicey. Suppose someone opts to pay a higher monthly fee and sign up for a one-month subscription. This user is presumably unsure about how much she will enjoy PlayStation Plus, so she is paying double the lowest monthly fee in exchange for the right to cancel her subscription if she doesn't enjoy the service all that much. If the customer selects that option, she will soon see this screen:



Ok. So customers who sign up for a one-month membership at \$10 per month will have that membership automatically renewed, at twice the monthly cost of customers who sign up for a 12-month membership. Presumably a tiny fraction of one-month subscribers prefer autorenewal at a high monthly rate. But never fear, as the figure above shows, Sony will let those customers opt out of automatic renewal, provided they click through . . . at least five screens – Settings, Account Managements, Account Information, Wallet, and Purchase Settings, where they will see a button that lets them toggle off autorenewal.<sup>24</sup> A user who neither writes down the precise opt-out instructions nor takes a digital photograph of the screen above will be lost at sea – the different steps a user must go through are far from intuitive.

<sup>24</sup> It's actually even more cumbersome. When one of the authors opted to turn off automatic renewal the author was required to re-log in to the system with a username and password, even though the author was already logged in.

A cynical observer might view Sony as furthering two objectives here. First, Sony knows that a number of their one-month subscribers will be auto-renewed at a high monthly rate, and that's a lucrative source of revenue for the company. Second, Sony knows that some of its customers will grasp immediately how difficult opting out of automatic renewal is, think it through a bit, and then press cancel. Presumably most will then sign up for the twelve-month subscription that Sony probably prefers, whose automatic renewal feature is less blatantly problematic. Either way, Sony comes out ahead.

When evaluating potential dark patterns, we need to be sure that we can differentiate between true positives and false positives. So in this instance we would want to know whether Sony's user interface is the product of an intentional design choice, an accident, or an external constraint. We will admit to a lack of hard data on this (in contrast to the remainder of this data-heavy paper) but in retrospect it seems clear that almost nobody who signs up for a one-month subscription at a high rate will also prefer for that subscription to autorenew. Where we see a user interface nudge consumers towards a selection that is likely to be unpopular with them but profitable for the company, there is reason to think a dark pattern may exist.<sup>25</sup> But perhaps Sony's programmers didn't think of that at the time. Alternatively, maybe letting people opt out of autorenewal for a PlayStation Plus subscription on one screen is inherently cumbersome for one reason or another. In this instance, we can more or less rule out the innocent explanations. Tellingly, once a customer signs up for autorenewal Sony will let them turn it off without navigating through five or more screens.



The initial set-up and very difficult process for opting out of autorenewal at the outset seems to be a conscious and intentional choice by Sony. If we examine what Sony is doing through the lens of existing taxonomies we can see that it is combining several tactics that have been identified as dark patterns.

---

<sup>25</sup> The connection between the majority sentiment among consumers and the identification of dark patterns is explored more explicitly in STRAHILEVITZ ET AL., *supra* note 13, at 44. In this paper's companion piece, we devote more time and experimental energy towards identifying the expectations and preferences that most consumers share. See Lior Jacob Strahilevitz & Jamie Luguri, *Consumertarian Default Rules*, \_\_ J. CONTEMP. PROBLEMS \_\_ (forthcoming 2020); see also Franklin G. Snyder & Ann M. Mirabito, *Consumer Preferences for Performance Defaults*, 6 MICH. BUS. & ENTREPRENEURIAL L. REV. 35 (2016) (reporting the results of survey research into consumer preferences in other sales contexts).



In this instance, Sony is combining a false hierarchy (the different sizes of the buttons on the initial screen), the bait and switch (the one-month subscription looks like it offers an easy way to decline the product after the user experiences it, but given user inertia it's often an indefinite subscription with a higher monthly rate), preselection (the default choice is good for the company and bad for most one-month subscription consumers), a roach motel (opting out of automatic renewal is far more difficult and time-consuming than keeping the automatic renewal); and forced continuity (many users will wind up paying for the service for a lengthy period of time despite their initial intent not to do so). These dark patterns are used in combination, seemingly in an effort to manipulate users into either a long-term subscription or an automatically renewing indefinite subscription at a high monthly rate.

To review, there are a variety of dark patterns that are designed to nudge consumers into contractual arrangements that they presumably would not otherwise prefer, and these techniques appear to be employed by a variety of different ecommerce firms, from start-up apps to well-capitalized platforms like Ticketmaster and Sony. Ticketmaster and Sony have a lot of smart people who work for them, so one assumes that they are doing what they are doing because it is good for the firms' bottom lines. But beyond that intuition we lack reliable information about the effectiveness of these dark patterns in nudging consumers to behave in ways that maximize firm profits. Turning to Part II of our paper, which is the heart of the project, we will now attempt to fill that gap in the literature. In order to do that, we created a classic "bait and switch" scenario with a large sample of Americans online.

## II. An Experimental Test of the Effectiveness of Dark Patterns

Let's suppose Amazon or Microsoft was interested in testing the effectiveness of dark patterns. It would be easy to do so using their existing platform. They have an ongoing relationship with millions of customers, and many of them have already stored their credit card information to enable one-click purchasing. So they could beta-test different dark patterns on subsets of their user-base, exploiting randomization, and then track purchases and revenue to see what works. The risks of customer / employee blowback or legal liability would be the main constraints on what they could do.

For academics seeking to test the efficacy of dark patterns, the challenge is much more significant. Academic researchers generally do not have established relationships with customers (students aside, and that relationship is heavily regulated where financial aid and tuition payment are concerned). The point of a dark pattern typically is to manipulate people to pay for something they otherwise would not purchase or surrender personal information they would otherwise keep confidential. There has been a little bit of academic work that has studied how different user interfaces can encourage the latter,<sup>26</sup> and none on the former. Because we are most interested in understanding how effective dark patterns are at parting consumers with their money, we wanted to situate ourselves in Amazon or Microsoft's shoes to the fullest extent possible. Alas, setting up a new ecommerce platform to run those experiments was prohibitively expensive.

---

<sup>26</sup> See, e.g., Laura Brandimarte et al., *Misplaced Confidences: Privacy and the Control Paradox*, 4 SOCIAL PSYCH. & PERSONALITY SCIENCE 340 (2012); Leslie K. John et al., *Strangers on a Plane: Context-Dependent Willingness to Disclose Sensitive Information*, 37 J. CONSUMER RES. 858 (2011); Marianne Junger et al., *Priming and Warnings Are Not Effective to Prevent Social Engineering Attacks*, 66 COMPUTERS IN HUM. BEHAV. 75 (2017).

To that end, we designed a bait-and-switch scenario that would strike consumers as plausible. We would use an existing survey research firm to recruit a large population of American adults to participate in a research study that would evaluate their attitudes about privacy. Then we would deceive those adults into believing, at the end of the survey, that because they expressed a strong interest in privacy (as respondents typically do in surveys) we had signed them up for a costly identity-theft protection service and would give them the opportunity to opt out. We would structure the experiment in a way so as to make experimental subjects believe that their own money was at stake and they would incur a legal obligation to pay for the service if they did not opt out, even though we did not have their credit card or bank payment information. Then we would randomly vary whether the opportunity to opt out was unconstrained or impeded by different dosages of dark patterns. This manipulation would plausibly generate information about consumers' revealed preferences, and it would allow us to do so without actually selling any goods or services to consumers. Our host university's I.R.B. approved our proposal to engage in the deceptive experiment after we explained, among other things, (a) that we wouldn't actually store any of the information that we were purportedly collecting to uniquely identify our experimental subjects, and (b) that we would promptly debrief participants at the end of the survey so they understood they would not be charged any money for our non-existent identity theft protection service.

To put that research plan in motion, we administered an online survey to a nationally representative (census weighted) sample of American participants recruited by Dynata, a professional survey research firm. We removed respondents who took too long or too little time to complete the survey from the sample, as well as those who failed an attention check.<sup>27</sup> This left a final sample of 1,963 participants.<sup>28</sup> Participants were compensated by Dynata for their time, and we compensated Dynata. We pre-registered the experiment with AsPredicted.Org.<sup>29</sup>

To begin, study participants answered various demographic questions including age, gender, race, education, income, employment, political orientation, and region of the country. Included with these basic demographic questions were additional questions aimed to bolster

---

<sup>27</sup> After removing two participants who started and ended the survey on different days, the average completion time was computed. Participants took 11.5 minutes on average to complete the survey. We removed participants who took less than 4 minutes and more than 47.5 minutes (two standard deviations above the survey completion time). Additionally, participants were asked an attention check question that asked them to "Please select "Strongly agree" for this question below to show that you are paying attention." Those that failed to answer accordingly were removed from the sample. At the end of the survey participants were asked to indicate how seriously they took the survey on a scale from 1 ("not at all seriously") to 5 ("extremely seriously"). Participants who answered 1 were also removed from the sample.

<sup>28</sup> Males comprised 47.1% of the sample. 76.2% of the sample self-identified as White, 13.2% as Black, 1.2% as American Indian, 4.4% as Asian, and 4.9% as "other." On a separate question, 14% indicated they are Hispanic or Latino. 6% of the sample had not completed high school, 29.8% had high school diplomas, 29.8% had some college or an associate's degree, 20.9% had bachelor's degrees, and 13.6% had advanced degrees. 10.8% of the sample was between 18-24 years old, 18% was between 25-34, 17.6% was between 35-44, 17.2% was between 45-54, 19.3% was between 55-64, and 17% was 65 years or older. 1773 participants (90.3%) fully completed the survey from start to finish.

<sup>29</sup> See [https://aspredicted.org/see\\_one.php](https://aspredicted.org/see_one.php) (Experiment # 19680) (submitted Feb. 17, 2019). On the value of pre-registration in social science research, see Brian A. Dosek et al., *The Preregistration Revolution*, 115 PNAS 2600 (2018).

the later cover story that we had pinpointed their mailing address. Specifically, participants were asked their zip code, how long they had lived at their current residence, their telephone number, and where they were completing the survey (home, work, or other).<sup>30</sup>

Participants then filled out the Ten Item Personality Measure, an instrument designed to measure the Big Five personality dimensions (extraversion, agreeableness, conscientiousness, emotional stability, and openness to experiences).<sup>31</sup> We included this measure to test whether certain personality traits would predict susceptibility to dark pattern manipulation.

Next, we assessed subjects' attitudes and opinions on data privacy. Though not the focus of the present paper, this section consisted of us asking participants about what companies either should be allowed to do or are allowed to do with consumer data. These questions focused on data collection, retention, third party sharing, and encryption. We collected responses to a battery of information privacy questions. This data collection allowed us to create a cover story for offering the participant identity theft protection.

Answering these questions took up most of respondents' time. In the last few minutes of the survey, they were exposed to a manipulation designed to assess the effectiveness of dark patterns. The first part of the survey provided us with an appropriate pretext for what followed. Respondents were told to wait while the software calculated their "privacy propensity score." All respondents were then told that based on their responses, our system had identified them as someone with a heightened concern about privacy. As such, we would automatically sign them up to receive a data and identity theft protection plan offered by our corporate partner, the largest and most experienced identity theft prevention and credit monitoring company in the United States. This was our bait and switch.

We told participants that by using the demographic information they had provided at the beginning of the survey, along with their IP address, we had pinpointed their mailing address. Our corporate partner would now provide them with six months of free data protection and credit history monitoring. After the six-month period, they would be billed monthly (though they could cancel at any time). The amount they would be billed varied by condition. Participants in the low stakes condition were told that the monthly fee would be \$2.99, and participants in the high stakes condition were told the fee would be \$8.99 per month.

Participants were then allowed to either accept or decline the data protection program. But the steps that were required to do so varied by the level of the dark pattern manipulation. In the control group condition, we did not include any dark patterns. As such, this condition serves as a baseline to help us establish a ceiling for what percentage of the sample was inherently interested in receiving the identity theft protection.<sup>32</sup> Participants could thus either

---

<sup>30</sup> In order to preserve confidentiality, these responses were deleted from the data set and were not analyzed.

<sup>31</sup> Samuel D. Gosling et al., *A Very Brief Measure of the Big Five Personality Dimensions*, 37 J. RES. PERS., 504, 525 (2003).

<sup>32</sup> We refer to this figure as a ceiling in the sense that it likely overestimates demand for the service subjects told our corporate partners were selling. This overestimation arises for at least two reasons. First, respondents were told that they already had been signed up for the service (potentially triggering loss aversion at the prospect of its removal) and second, subjects were told that they would pay nothing for the service for the first six months (potentially triggering hyperbolic discounting and optimism bias about whether their future selves would remember to cancel the service once the free trial period ended). We had also primed them to think a lot about privacy, though it is not clear which way that cut, given our

click “Accept” or “Decline” on the first screen. Regardless of which option they selected, they proceeded to the final stage of the experiment, which is described below.

In the mild dark patterns condition, subjects could either click “Accept and continue (recommended) or “Other options,” and the button that accepted the program was selected by default. We made it easier for users to accept the program (because they did not have to select the button themselves) and harder to decline it (because there was not a straightforward and immediate way to decline, only to see other options). Adding a “recommended” parenthetical is a form of false hierarchy. The parenthetical plausibly triggers a heuristic where consumers encounter recommendations made by a neutral fiduciary elsewhere and may be uncertain as to who is making the recommendation and what the basis for that recommendation is.<sup>33</sup>

If subjects selected “Other options,” they were directed to the next screen, which asked them to choose between “I do not want to protect my data or credit history” or “After reviewing my options, I would like to protect my privacy and receive data protection and credit history monitoring.” This question uses confirmshaming as a dark pattern to nudge respondents to accept the program (i.e. their decision to decline the program is framed as not wanting to protect their data).

Next, if subjects did not accept the program, they were asked to tell us why they declined the valuable protection. Several non-compelling options were listed, including “My credit rating is already bad,” “Even though 16.7 million Americans were victimized by identity theft last year, I do not believe it could happen to me or my family,” “I’m already paying for identity theft and credit monitoring services,” and “I’ve got nothing to hide so if hackers gain access to my data I won’t be harmed.” They also could choose “Other” and type in their reason, or choose “On second thought, please sign me up for 6 months of free credit history monitoring and data protection services.” This is another confirmshaming strategy. Additionally, it makes it more onerous for many users to decline rather than accept (because if they did not select one of the sub-optimal options provided, they were asked to type out their reason for declining). Subjects who rejected the data protection plan on this screen were treated as having declined the service, and they advanced to the same final screens that those in the control group also saw.

In the aggressive dark pattern condition, the first and second screens were identical to those in the mild dark pattern condition. Participants attempting to decline the identity theft protection were then told that since they indicated they did not want to protect their data, we would like to give them more information so they could make an informed choice. We asked them to read a paragraph of information about what identity theft is. Participants could either choose “Accept data protection plan and continue” or “I would like to read more information.” They were forced to remain on the page for at least ten seconds before being able to advance, and they were shown a countdown timer during this period. This screen created a significant roach motel. Namely, it obstructed respondents’ ability to decline the program by making it more onerous to decline than accept.<sup>34</sup> It also toyed with respondents’ emotions by using vivid,

---

setup. Because we are more interested in comparing the control group to the dark pattern conditions than we are in estimating the actual unmet demand for an identity theft protection service, this potential overestimation presents no problems for our study.

<sup>33</sup> Gray et al., *supra* note 4, at 7.

<sup>34</sup> *Id.* at 6.

frightening language in the text. For example, participants read that identity theft “can damage your credit status, and cost you time and money to restore your good name.”

If respondents chose to read more information (rather than accept the program), the next screen had information about why identity theft matters and what a thief could do with their personal information. The options and countdown timer were the same as the previous screen. A third information screen explained how common identity theft is, with the same options and countdown timer displayed before they could advance. The cumulative effect of these screens amounted to a nagging dark pattern.

If participants endured all three information screens and chose “I would like to read more information,” they were then directed to a question designed to confuse them. They were asked, “If you decline this free service, our corporate partner won’t be able to help you protect your data. You will not receive identity theft protection, and you could become one of the millions of Americans who were victimized by identity theft last year. Are you sure you want to decline this free identity theft protection?” The two options were “No, cancel” and “Yes.” This trick question intentionally tried to confuse participants about which option they should select to decline the program.<sup>35</sup> Checking the box that includes the word “cancel” counterintuitively accepts the identity theft program. Participants choosing “Yes” were directed to the same last screen as in the mild dark pattern condition, which asked them to indicate their reason for declining the program. After that, they were sent to the same final screens that all subjects saw.

At the conclusion of the study, all participants were asked to indicate their current mood.<sup>36</sup> They were then asked whether they would be interested in potentially participating in follow-up research studies by the same researchers.<sup>37</sup> Next, they were asked how free they felt they were to refuse the offered plan.<sup>38</sup> These questions aimed to assess whether companies that employ dark patterns face any negative repercussions for their use. By comparing the responses of mild and aggressive dark pattern participants to those of the control group we could estimate the size of the good-will loss that a company employing dark patterns would suffer. Lastly, participants were asked how seriously they took the survey, and then were given a text box to write any questions, comments, or concerns they had about the survey. They were then thoroughly debriefed.

#### A. Rates of Acceptance

The results of the study offer striking empirical support for the proposition that dark patterns are effective in bending consumers’ will. As expected, in the control group condition, respondents opted to accept the identity theft protection program at very low rates. Only 11.3% of respondents accepted the program when they were allowed to accept or decline the program

---

<sup>35</sup> For a discussion of similar dark pattern strategies in Apple’s iOS 6, see WOODROW HARTZOG, *PRIVACY’S BLUEPRINT* 208 (2018).

<sup>36</sup> Participants indicated their mood on a scale from 1 (“Happy and relaxed”) to 7 (“Aggravated and annoyed”).

<sup>37</sup> Participants indicated their interest on a scale from 1 (“Not at all”) to 7 (“Extremely interested”).

<sup>38</sup> Participants indicated their degree of freedom on a scale from 1 (“Not at all free to refuse”) to 7 (“completely free to refuse”).

on the first screen.<sup>39</sup> This acceptance rate likely overestimates the demand for a product of this kind.<sup>40</sup>

When mild dark pattern tactics were deployed, the acceptance rate more than doubled. Now 25.8% of participants accepted the data protection program, which corresponds to a 228% increase compared to the control group condition. When participants were exposed to aggressive dark patterns aggressive, the acceptance rate shot up further, with 41.9% of the sample accepting the program.<sup>41</sup> So the aggressive dark pattern condition nearly quadrupled the rate of acceptance, with a 371% increase in rates of acceptance compared to the control group condition. These results are statistically significant and then some. The effect sizes are enormous by the standards of social science research. Manipulative tactics widely employed in the world of brick-and-mortar sales are evidently much less powerful in comparison.<sup>42</sup>

Table 3: Acceptance Rates by Condition

Condition	Acceptance Rate (%)	Number of Respondents Accepting
Control group	11.3%	73 (out of 644)
Mild	25.8%	155 (out of 600)
Aggressive	41.9%	217 (out of 518)

Given the experimental design, it is possible to determine when participants chose to accept the program in the mild and aggressive dark-pattern conditions. In other words, which of the dark pattern questions seemed to be doing the “work” in nudging participants toward accepting the program? In both conditions, the initial screen (which offered a choice between “Accept and continue (recommended)” and “Other options,” with the former choice pre-selected) accounted for the majority of acceptances. In the mild condition, more than three-quarters of participants who accepted did so on this first screen (75.5%, 117 out of 155). In the aggressive condition, this screen accounted for 65% of acceptances (141 out of 217).<sup>43</sup> The

---

<sup>39</sup> Participants were counted as accepting the program if, in any question, they selected the option to accept. They were counted as declining if they indicated they declined in the control group condition, or if they reached the last screen in the mild and aggressive conditions and selected an option other than accepting the program. Therefore, participants who dropped out during the dark pattern manipulation were neither counted as accepting or declining the program. In Table 3 above, if we were to count those who dropped out as decliners, the acceptance rate for the mild dark patterns group would fall to 25% and the acceptance rate for the aggressive dark patterns group would fall to 37%.

<sup>40</sup> See *supra* note 32.

<sup>41</sup> A chi-square test of independence was performed to examine the relationship between dark pattern condition and acceptance rates. The relation between these variables was significant,  $\chi^2(2, N=1762)=142.16, p<.001$ .

<sup>42</sup> See, e.g., Hanson & Kysar, *supra* note 3, at 1447-48 (noting that clever strategies designed to increase impulse buying had boosted sales by “at least ten percent” and placing products at eye-level rather than on a low shelf at a grocery store increases toothbrush purchases by 8 percent).

<sup>43</sup> Because the aggressive dark pattern subjects had more opportunities to relent and accept the data protection plan later in the survey it makes sense that this percentage is lower. The higher dropout rate of those in the aggressive dark patterns condition, discussed below, is another contributing factor. Counting

second screen (which offered a choice between “I do not want to protect my data or credit history” and “After reviewing my options, I would like to protect my privacy and receive data protection and credit history monitoring”) accounted for 35 more acceptances in the mild condition (23% of overall acceptances) and 22 more in the aggressive condition (10% of acceptances). For those in the aggressive condition, when participants were forced to read three screens of information on identity theft for at least ten seconds per screen, this roach motel and nagging strategy accounted for 19% of acceptances overall. The confusing trick question (offering an “Are you sure you want to decline this free identity theft protection?” prompt with the options “No, cancel” and “Yes.”) was responsible for another 11% of acceptances. Finally, nearly no one who made it to the final, confirmshaming screen (the list of largely bad reasons for declining the service, with one final chance to accept) ended up accepting, either in the mild or aggressive conditions.

Of course, participants only advanced to new screens in the mild and aggressive conditions if they didn’t “fall” for the dark pattern on an earlier screen. As soon as they accepted the program, the dark patterns ceased (as is often the case in the real world). This means that it is not correct to infer the relative strengths of the different dark patterns deployed from the number of acceptances each caused. Dark patterns that were used later in the manipulation are less likely to work by the very fact that people who were most susceptible to dark patterns were no longer in the sample because they already accepted the data protection plan.

That said, the information about when people accepted is informative for two reasons. First, it demonstrates the substantial cumulative power that different kinds of dark patterns can have. Some people who were able to resist certain dark patterns (like roach motels) are still susceptible to falling for others (like confusingly worded questions). Second, this data demonstrates that seemingly minor dark patterns can have relatively large effects on consumer choices. In the control group condition, participants were able to choose “Accept” or “Decline.” Changing these options to “Accept and continue (recommended)” and “Other options,” with the former pre-selected, all by itself, nearly doubled the percentage of respondents accepting the program.<sup>44</sup>

## B. The Influence of Stakes

Across the dark pattern conditions, we varied the price point of the program (\$2.99 vs. \$8.99) to see whether higher monetary stakes influenced rates of acceptance. The neoclassical model of economics generally predicts that consumers will be willing to jump over more hurdles in order to save themselves more money. On this account, consumers face a tradeoff between out-of-pocket expenses to be incurred later and annoying wasted time costs to be incurred now. Impatient consumers should therefore be more likely to relent and accept the program when

---

dropouts as people who declined the data protection plan, 19.2% of subjects in the mild dark pattern condition and 24.2% of subjects in the aggressive dark pattern condition accepted the offer on the first screen. For a full breakdown of acceptance rate by question, see Appendix A.

Our intuition about this data is that a number of consumers have encountered dark patterns in the wild before and they feel that they can surrender to them now or surrender to them later, so they may as well surrender early and save themselves some time. That said, further studies that randomize the order of dark patterns would be necessary to confirm or refute this hypothesis.

<sup>44</sup> The acceptance rate increased from 11.3% to 20.7% in the combined dark patterns conditions (counting only those users who accepted on the first dark pattern screen).

the costs of acceptance are lower.<sup>45</sup> Moreover, rational consumers should be more attentive on average when they are asked to pay a higher price for a good or service, and this might make them less prone to mistakes or impulsive decision-making.

On the basis of these neoclassical assumptions one of the authors (who has produced a fair bit of law & economics scholarship) hypothesized that in the high-stakes condition, overall acceptance rates would be lower. Additionally, he predicted that when respondents had more money at stake, they would be less likely to “fall” for the dark patterns employed in the mild and aggressive conditions. The other author (a psychologist) expressed her consistent skepticism about this hypothesis. The predictions suggested by the neo-classical model were not borne out by the data, and the psychologist’s skepticism proved well-founded. Rates of acceptance were not related to stakes.<sup>46</sup> There were no significant differences between the high- and low-stakes conditions across any of the dark pattern conditions (see Appendix B for acceptance rates broken down by stakes and level of dark pattern).<sup>47</sup> Tripling the cost of a service had no effect on uptake in this domain. You read that right.<sup>48</sup>

### C. Potential Repercussions of Deploying Dark Patterns

The rates of acceptance in the mild and aggressive conditions show that dark patterns are effective at swaying consumer choices. Though only a small percentage of participants were truly interested in the data protection program for its own sake, a much larger percentage decided to accept the program after we exposed them to dark patterns. These results illustrate why dark patterns are becoming more common — because companies know that they are effective in nudging consumers to act against their own preferences. But it is possible that companies experience a backlash by consumers when they use dark patterns. If so, then there would be less concern that dark patterns are the result of market failure, weakening the case for legal intervention. The questions asked immediately after the experiment were designed to get at this question.

First, participants were asked about their mood to assess whether exposure to dark patterns elicited negative emotions. There was an overall effect of the dark pattern manipulation.<sup>49</sup> While participants in the control group ( $M=2.96$ ,  $SD=1.61$ ) and mild ( $M=3.05$ ,

---

<sup>45</sup> One countervailing force consistent with the neoclassical model is that high price can function as a signal of quality. See, e.g., Ayelet Gneezy et al., *A Reference-Dependent Model of the Price-Quality Heuristic*, 51 J. MARKETING RES. 153, 154 (2014). There is obviously a limit to this signaling dynamic, however, which constrains price increases.

<sup>46</sup> A chi-square test of independence was performed to examine the relationship between stakes (low vs. high) and acceptance rates. The relation between these variables was not significant,  $\chi(1, N=1762)=0.76$ ,  $p=.38$ .

<sup>47</sup> Chi-square tests for independence were run separately for each of the dark pattern conditions. There was no significant relationship between stakes and acceptance rates in the control group ( $\chi(1, N=644)=2.52$ ,  $p=.11$ ), mild ( $\chi(1, N=600)=0.27$ ,  $p=.61$ ), or aggressive ( $\chi(1, N=518)=0.19$ ,  $p=.66$ ) conditions.

<sup>48</sup> One possible explanation for these results is that consumers in the high-stakes condition felt they were getting six months of very valuable data protection for free, whereas those in the low-stakes condition felt they were getting six months of less valuable data protection for free. It is possible that the greater perceived upside of the six month free trial cancelled out the greater perceived downside of paying \$8.99 per month once the trial period ended.

<sup>49</sup>  $F(2,1740)=323.89$ ,  $p<.001$



SD=1.73) conditions reported similar levels of negative affect, participants in the aggressive condition were significantly more upset ( $M=3.94$ ,  $SD=2.06$ ).<sup>50</sup> These results suggest that if companies go too far and present customers with a slew of blatant dark patterns designed to nudge them, they might experience backlash and the loss of good will. Yet it is notable that the mild dark pattern condition more than doubled the acceptance rate and did not prompt discernable emotional backlash.

At the end of the study, participants had another chance to express their emotions; they were given a box to type any questions, concerns, or comments they might have. We decided to code these responses to see whether, similar to the explicit mood question mentioned above, participants were more likely to spontaneously express anger after having been exposed to the mild or aggressive dark patterns.<sup>51</sup> The pattern of results mirrored those of the explicit mood measure. Participants in the control group and mild conditions did not express anger at different rates. However, participants in the aggressive dark pattern condition were significantly more likely to express anger.<sup>52</sup>

Taken together, these two mood measures suggest that overexposure to dark patterns can irritate people. Respondents in the aggressive dark pattern condition reported being more aggravated, and were more likely to express anger spontaneously. It is notable that those respondents exposed to the mild dark patterns did not show this same affective response. Though the mild condition very substantially increased the percentage of respondents accepting the data protection program, there were no corresponding negative mood repercussions.

Even though respondents in the aggressive dark pattern condition overall expressed more negative affect (thereby indicating a potential backlash) it is important to understand what is driving this aggravation. Are people who end up accepting or declining the program equally angered by the use of dark patterns? To answer this question, we compared the moods of people who accepted or declined across the dark pattern conditions. There was an overall main effect, such that people who declined the program reported more displeasure ( $M=3.50$ ,  $SD=1.99$ ) than those who accepted the program ( $M=3.21$ ,  $SD=1.78$ ).<sup>53</sup> This effect is driven by the aggressive dark pattern condition. Specifically, among people who accepted the program, there was no significant differences in mood across the control group, mild, and aggressive dark

---

<sup>50</sup> Post-hoc Tukey HSD tests confirmed that mean differences in the control group and mild conditions were not significant ( $p=.63$ ), but both differed significantly from the aggressive condition ( $p<.001$ ).

<sup>51</sup> Participants who did not write anything, wrote something neutral, or wrote something positive were coded as a 0. Participants who either expressed general anger or anger specifically at the offer of the data protection program were coded as a 1.

<sup>52</sup> In the control group condition, 36 out of 632 (5.70%) were coded as expressing anger. In the mild condition, the rate was 36 out of 591 (6.09%). In the aggressive condition, it was 66 out of 515 (12.82%). A chi-square test of independence was performed to examine the relationship between dark pattern condition and whether anger was expressed (Yes/No). The relation between these variables was significant,  $\chi(1, N=1738)=23.86$ ,  $p<.001$ . The control group and mild conditions did not differ significantly from each other ( $\chi(1, N=1151)=0.09$ ,  $p<.77$ ) but both differed significantly from the aggressive condition (control group vs. aggressive:  $\chi(1, N=1045)=17.75$ ,  $p<.001$ ; mild vs. aggressive:  $\chi(1, N=1004)=14.86$ ,  $p<.001$ ).

<sup>53</sup>  $F(1,1741)=8.21$ ,  $p=.004$ .

pattern conditions.<sup>54</sup> However, among those who declined, respondents in the aggressive dark pattern condition were more aggravated than those in the control group and mild conditions. The latter two conditions did not differ. This suggests that when dark patterns are effective at leading people to a certain answer, there is no affective backlash. Only when participants are forced to resist a slew of dark patterns in order to express their preference do we observe increased aggravation.

In addition to mood, another potential kind of backlash that dark patterns might elicit is disengagement. People might negatively react because they feel pressured, leading them to want to avoid the dark patterns either in the moment or be hesitant to interact with the entity that employed the dark patterns in the future. In the current study, we have two measures that capture this potential disengagement.

First, participants were able to exit the survey at any time, though if they failed to complete the survey they forfeited the compensation to which they'd otherwise be entitled. We therefore can examine whether participants were more likely to drop out of the study in the aggressive versus mild conditions.<sup>55</sup> We found that respondents were much more likely to drop out and disengage with the study in the aggressive condition.<sup>56</sup> Only 9 participants dropped out in the mild condition, while 65 dropped out at some point during the aggressive condition. The latter is an unusual, strikingly high dropout rate in our experience, made all the more meaningful by the sunk costs fallacy. Respondents had typically devoted ten minutes or more to the survey before encountering the dark pattern, and by exiting the survey during the dark pattern portion of the experiment they were forfeiting money they may well have felt like they had already earned.<sup>57</sup>

Second, participants were told that some of them might be contacted to do a follow up survey with the same researchers. They were asked if they were potentially interested in participating. We expected participants to be less interested in the follow-up study if they had been exposed to the mild or aggressive dark pattern conditions. The results supported this

---

<sup>54</sup> There was a significant interaction between dark pattern manipulation and outcome,  $F(5,1737)=15.12$ ,  $p<.001$ . Among people who accepted, there was no main effect of dark pattern condition,  $F(2,434)=0.62$ ,  $p=.54$ . However, among those who declined, there was a main effect,  $F(2,1303)=67.02$ ,  $p<.001$ . Post-hoc Tukey tests revealed that respondents who declined after being exposed to the aggressive dark pattern condition were significantly more aggravated than those in the control group and mild conditions ( $ps<.001$ ). Respondents who declined in the control group and mild conditions did not differ significantly,  $p=.81$ .

<sup>55</sup> Because the control group condition only contained one question, there was no opportunity for participants to drop out in this condition.

<sup>56</sup> A chi-square test of independence was performed to examine the relationship between dark pattern condition (mild vs. aggressive) and whether participants dropped out or not. The relation between these variables was significant,  $\chi(1, N=1192)=47.85$ ,  $p<.001$ .

<sup>57</sup> The dropout rates observed provide highly relevant information about the social welfare costs of dark patterns. A reasonably high percentage of respondents were willing to forfeit real money rather than continuing to incur the costs of declining an unwanted service or running the risk that they would be signed up for a service they did not want. Of course, by closing their browser and stopping the experiment, there was no guarantee that they would avoid the unwanted subscription. We told respondents at the beginning of the experiment that we had already signed them up for the data protection plan using information they had provided at the beginning of the survey.

hypothesis. Dark pattern condition was significantly related to interest in participating in a follow-up survey.<sup>58</sup> Participants in the control group condition indicated significantly more interest ( $M=4.46$ ,  $SD=2.31$ ) than participants in the mild ( $M=4.11$ ,  $SD=2.32$ ) and aggressive ( $M=3.97$ ,  $SD=2.39$ ) conditions.<sup>59</sup> However, here the difference between those in the mild and aggressive conditions was not significant.<sup>60</sup> This is the one measure of customer sentiment where significant differences were observed between the control group and subjects exposed to mild dark patterns.

One potential reason for the disengagement found above is that the more participants were exposed to dark patterns, the more likely they were to feel coerced into accepting the data protection program. To assess this, we asked participants how free they felt to refuse the data protection program. As expected, condition significantly influenced feelings of freedom.<sup>61</sup> Participants in the control group condition felt freer to refuse ( $M=6.21$ ,  $SD=1.44$ ) compared to those in the mild ( $M=5.81$ ,  $SD=1.75$ ) and aggressive ( $M=4.74$ ,  $SD=2.26$ ) conditions.<sup>62</sup> Interestingly, as the median scores suggest, most respondents felt more free than unfree to refuse the program, even in the aggressive dark pattern condition.

#### D. Predicting Dark Pattern Susceptibility

Given the strong influence that dark patterns seem to exert on consumer choice, it is important to understand what individual differences might predict susceptibility. Put another way, what kinds of people are more vulnerable to being manipulated by dark patterns?<sup>63</sup> To answer this question, we analyzed whether demographic and personality differences predicted acceptance rates across dark pattern conditions.

We first analyzed whether education predicts acceptance of the program and found that it does.<sup>64</sup> The less educated participants were, the more likely they were to accept the program. The key question, though, is whether the relationship between level of education and likelihood of acceptance varies by dark pattern condition. In the control group condition, education is not significantly related to whether participants accepted or declined.<sup>65</sup> This means that in the absence of dark patterns, participants with high and low levels of education do not differentially value the offered program. However, when they are exposed to mild dark patterns, participants

---

<sup>58</sup>  $F(2,1740)=6.99$ ,  $p=.001$ .

<sup>59</sup> Post-hoc Tukey tests revealed that participants in the control group condition differed significantly from both those in the mild ( $p=.02$ ) and aggressive ( $p=.001$ ) conditions.

<sup>60</sup>  $p=.57$ .

<sup>61</sup>  $F(2,1739)=96.63$ ,  $p<.001$ .

<sup>62</sup> Post-hoc Tukey tests reveal that all three conditions are significantly different from one another,  $p_s<.001$ .

<sup>63</sup> In other contexts, scholars have found that people with fewer financial resources have more difficulty overcoming administrative burdens that people with more resources. See PAMELA HERD & DANIEL P. MOYNIHAN, ADMINISTRATIVE BURDEN: POLICYMAKING BY OTHER MEANS 7-8, 57-60 (2019).

<sup>64</sup> A logistic regression was performed to ascertain the effects of education on the likelihood that participants accepted the data protection program. Level of education significantly predicted whether participants accepted or declined the program,  $b=-.15$ ,  $SE=.04$ ,  $p<.001$ , such that participants with greater levels of education were more likely to decline.

<sup>65</sup>  $b=-.11$ ,  $SE=.08$ ,  $p=.17$ .

with less education become significantly more likely to accept the program.<sup>66</sup> A similar pattern of results emerged in the aggressive dark pattern condition.<sup>67</sup>

When controlling for income, the relationship between education and acceptance varies slightly. The results are similar, except that less education no longer predicts acceptance in the aggressive dark pattern condition.<sup>68</sup> The relationship persists in the mild dark pattern condition with these controls. This pattern of results endures when additional demographic variables are controlled for.<sup>69</sup> This result further illustrates the insidiousness of relatively mild dark patterns. They are effective, engender little or no backlash, and exert a stronger influence on more vulnerable populations.

Next, we examined whether political ideology predicted acceptance across dark pattern conditions. Mirroring the results of education, in the control group condition political ideology does not predict acceptance.<sup>70</sup> But in the mild and aggressive conditions, participants who were more conservative were more likely to accept.<sup>71</sup> This pattern of results remains even when demographic differences are controlled for.<sup>72</sup> The results are interesting, though the effect sizes are not especially large.

Lastly, we examined whether personality traits predicted susceptibility to dark patterns. At the beginning of the survey, participants filled out a personality inventory that measured the Big 5 traits: extraversion, agreeableness, conscientiousness, neuroticism, and openness. Looking across dark pattern conditions, only extraversion and conscientiousness predict acceptance (See Appendix C for full analyses of all five personality traits).<sup>73</sup> More extraverted people and less conscientious people are more likely to accept the program. Breaking down these results by dark pattern condition, the relationship between extraversion and conscientiousness remain in the control group and mild conditions.<sup>74</sup> However, both traits fail to predict behavior (accepting or declining the program) in the aggressive condition. This result is particularly notable, and confusing, for conscientiousness. People who are conscientious tend to be more diligent and careful. One might expect this personality trait to offer insulation from the manipulative effects of dark patterns. Yet when participants are exposed to a slew of dark patterns in the aggressive

---

<sup>66</sup>  $b = -.19$ ,  $SE = .06$ ,  $p = .002$ .

<sup>67</sup>  $b = -.17$ ,  $SE = .06$ ,  $p = .003$ .

<sup>68</sup> Control group condition:  $b = -.08$ ,  $SE = .09$ ,  $p = .40$ . Mild condition:  $b = -.17$ ,  $SE = .07$ ,  $p = .01$ . Aggressive condition:  $b = -.07$ ,  $SE = .07$ ,  $p = .27$ .

<sup>69</sup> Controls include income, age, gender, and race (white vs. non-white). Control group condition:  $b = -.05$ ,  $SE = .10$ ,  $p = .57$ . Mild condition:  $b = -.18$ ,  $SE = .07$ ,  $p = .01$ . Aggressive condition:  $b = -.08$ ,  $SE = .07$ ,  $p = .24$ .

<sup>70</sup>  $b = .00$ ,  $SE = .07$ ,  $p = 1.0$ .

<sup>71</sup> Mild condition:  $b = .12$ ,  $SE = .06$ ,  $p = .03$ . Aggressive condition:  $b = .13$ ,  $SE = .05$ ,  $p = .01$ .

<sup>72</sup> Controls include gender, age, education, income, and race (white vs. non-white). Control group condition:  $b = .02$ ,  $SE = .07$ ,  $p = .79$ . Mild condition:  $b = .13$ ,  $SE = .06$ ,  $p = .03$ . Aggressive condition:  $b = .15$ ,  $SE = .05$ ,  $p = .007$ .

<sup>73</sup> Logistic regressions were run controlling for education, income, gender, age, and race to examine the relationship between extraversion ( $b = .10$ ,  $SE = .04$ ,  $p = .02$ ) and conscientiousness ( $b = -.16$ ,  $SE = .05$ ,  $p = .001$ ) on acceptance rates.

<sup>74</sup> Extraversion only marginally predicts acceptance in the mild condition. See Appendix C.

condition, we do not see different acceptance rates among those who are more or less conscientious.

To summarize the data we have collected and analyzed here, it appears that dark patterns can be very effective in prompting consumers to select terms that substantially benefit firms. These dark patterns might involve getting consumers to sign up for expensive goods or services they do not particularly want, as in our study and several real-world examples discussed in the previous part, or they might involve efforts to get consumers to surrender personal information – a phenomenon we did not test but that also is prevalent in ecommerce.

From our perspective, it's the mild dark patterns tested – like labeling an option that is good for a company's bottom line but maybe not for consumers as "recommended" or by providing initial choices between "Yes" and "Not Now" – that are most insidious. This kind of decision architecture, combined with the burden of clicking through an additional screen, managed to more than double the percentage of respondents who agreed to accept a data protection plan of dubious value, and it did so without alienating customers in the process. As a result, consumers were manipulated into signing up for a service that they probably did not want and almost certainly did not need. More broadly, we can say the same things about the kinds of dark patterns that are proliferating on digital platforms. These techniques are harming consumers by convincing them to surrender cash or personal data in deals that do not reflect consumers' actual preferences and may not serve their interests. There appears to be a substantial market failure where dark patterns are concerned – what is good for ecommerce profits is bad for consumers, and plausibly the economy as a whole. Legal intervention is justified.<sup>75</sup>

We now know that dark patterns are becoming prevalent and they can be powerful. Knowing these things raises the question of whether they are also unlawful (as unfair or deceptive practices in trade). It also implicates the related question of whether consumer assent secured via dark pattern manipulations ought to be regarded as consent by contract law. Finally, if readers conclude that dark patterns ought to be unlawful or ought not to count as valid consumer consent, that conclusion raises a host of implementation issues. Front and center, can the legal system draw stable lines between permissible (and constitutionally protected) commercial persuasion and impermissible dark patterns? We consider those issues in Part III.

### III. Are Dark Patterns Unlawful?

There are several plausible legal hooks that could be used to curtail the use of dark patterns in ecommerce. First, the Federal Trade Commission Act restricts the use of unfair or deceptive practices in interstate trade, providing the Commission with a mandate to regulate and restrict such conduct. Second, state unfair competition laws include similar frameworks. Finally, there is a broad question about whether consumer consent that is procured in a process that employs highly effective dark patterns should be voidable, which would entitle consumers to various remedies available under contract law and which could open up liability for firms that engage in various activities (for example, engaging in surveillance or processing biometric information) without having first obtained appropriate consumer consent.

---

<sup>75</sup> See Sunstein, *Sludge Audits*, *supra* note 2, at 15.

### A. Laws Governing Deceptive and Unfair Practices in Trade

The F.T.C., with its power to combat unfair and deceptive acts and practices under section 5 of the F.T.C. Act, is the most obvious existing institution that can regulate dark patterns. The scope of the F.T.C.'s investigation and enforcement authority covers "any person, partnership or corporation engaged in or whose business affects commerce,"<sup>76</sup> with some minor exceptions. As such the F.T.C. has the necessary reach to restrict the use of dark patterns across a wide range of industries. Since 1938 the F.T.C. Act has included language prohibiting "unfair or deceptive acts or practices in or affecting commerce."<sup>77</sup> The scope of the F.T.C.'s reach and the language of the provision remains broad, reflecting Congress's view that it would be challenging to specify *ex ante* all the different forms of behavior in trade that might be problematic. The Judiciary has consistently deferred to the F.T.C.'s interpretation of its mandate, with the Supreme Court holding in *F.T.C. v. Sperry & Hutchinson Co.*, that the F.T.C. Act allows, "the Commission to define and proscribe practices as unfair or deceptive in their effect upon consumers."<sup>78</sup>

In using its authority to restrict deceptive acts or practices affecting commerce, the F.T.C. treats as deceptive any "representation, omission, or practice" that is (a) material, and (b) likely to mislead consumers who are acting reasonably under the circumstances.<sup>79</sup> Materiality involves whether information presented "is important to consumers and, hence, likely to affect their choice of, or conduct regarding, a product."<sup>80</sup> Any express product claims made by a company are presumptively material.<sup>81</sup> As for the second prong, "the Commission need not find that all, or even a majority, of consumers found a claim implied" a false or misleading statement. Rather, liability "may be imposed if at least a significant minority of reasonable consumers would be likely to take away the misleading claim."<sup>82</sup> When enforcing the law, the F.T.C. need not show that the defendants intended to deceive consumers. Rather, it will be adequate for the agency to show that the "overall net impression" of the defendant's communication is misleading.<sup>83</sup> Thus, a company cannot make an initial series of misstatements and then bury the corrections of those misstatements in a subsequent communication.<sup>84</sup>

Because lawyers have written very little about dark patterns, and because computer scientists writing in the field are largely unaware of developments in the case law, the existing literature has missed the emergence in recent years of numerous F.T.C. enforcement actions that target dark patterns, albeit without using that term. Indeed, many of the key published

---

<sup>76</sup> 15 U.S.C. § 46(a).

<sup>77</sup> Matthew Sawchak & Kip Nelson, *Defining Unfairness in "Unfair Trade Practices"*, 90 N.C.L. REV. 2033 (2012).

<sup>78</sup> 405 U.S. 233 (1972).

<sup>79</sup> In the Matter of Cliffdale Assoc., Inc., 103 F.T.C. 110, 1984 WL 565319 (F.T.C. Mar. 23, 1984).

<sup>80</sup> *F.T.C. v. Cyberspace.com LLC*, 453 F.3d 1196, 1201 (9th Cir. 2006).

<sup>81</sup> *F.T.C. v. Pantron 1 Corp.*, 33 F.3d 1088 (9th Cir. 1994).

<sup>82</sup> *Fanning v. F.T.C.*, 821 F.3d 164, 170-71 (1st Cir. 2016).

<sup>83</sup> *F.T.C. v. E.M.A. Nationwide, Inc.*, 767 F.3d 611, 631 (6th Cir. 2014).

<sup>84</sup> *Id.* at 633.

opinions postdate Ryan Calo's survey of the law from 2014, in which he found hardly any relevant F.T.C. enforcement actions.<sup>85</sup>

*Federal Trade Commission v. AMG Capital Management* is the most important of the dark patterns cases, but because it's very recent and flew under the radar when it was decided it has not yet been discussed at all in the legal scholarship.<sup>86</sup> The dispute involved the F.T.C.'s enforcement action against a payday lender that was using various dodgy tactics to lure customers. The primary defendant, Scott Tucker, ran a series of companies that originated more than \$5 million in payday loans, typically for amounts less than \$1000.<sup>87</sup> Tucker's websites included Truth in Lending Act (TILA) statements explaining that customers would be charged a finance rate of, say, 30% for these loans. But the fine print below the TILA disclosures mentioned an important caveat. Amidst "densely packed text" especially diligent readers were informed that customers could choose between two repayment options – a "decline to renew" option and a "renewal" option.<sup>88</sup> Customers who wanted to decline to renew would pay off the payday loan at the first opportunity, provided they gave Tucker's company notice of their intention to do so at least three business days before the loan was due.<sup>89</sup> On the other hand, customers who opted for "renewal" would accrue additional finance charges, such as an additional 30 percent premium on the loan. After three such renewals, Tucker would impose an additional \$50 per month penalty on top of the accumulated premiums. As the Ninth Circuit explained, a typical customer who opted for the renewal option could expect to pay more than twice as much for the loan as a typical "decline to renew" customer.<sup>90</sup> So, of course, Tucker's companies made "renewal" the default option and buried information about how to switch to the "decline to renew" option in a wall of text.<sup>91</sup> That was the case even though the TILA disclosures provided the repayment terms under the assumption that a customer opted to decline to renew.

Judge O'Scannlain, writing for the court, was not impressed with Tucker's protestations that his disclosures were "technically correct." In the Court's view, "the F.T.C. Act's consumer-friendly standard does not require only technical accuracy.... Consumers acting reasonably under the circumstances – here, by looking to the terms of the Loan Note to understand their obligations – likely could be deceived by the representations made here. Therefore, we agree with the Commission that the Loan Note was deceptive."<sup>92</sup> Tucker's web sites employed numerous dark patterns. Renewal option customers were subjected to forced continuity (a costly subscription by default) and a roach motel (avoiding the onerous default is more taxing than submitting to it). And all customers had to overcome hidden costs (the burial of the renewal option's onerous terms in a long wall of text), preselection (making renewal the default), and trick question text (hard-to-understand descriptions of their options) in order to

---

<sup>85</sup> Calo, *supra* note 3, at 1002.

<sup>86</sup> 910 F.3d 417 (9th Cir. 2018).

<sup>87</sup> *Id.* at 420.

<sup>88</sup> *Id.* at 422.

<sup>89</sup> *Id.* at 423.

<sup>90</sup> *Id.*

<sup>91</sup> *Id.*

<sup>92</sup> *Id.* at 424.

avoid paying substantially higher fees. Each of these problematic aspects of the web site design was emphasized by the circuit court.<sup>93</sup> The court did not need a dark patterns label or experimental data to see how deceptive the individual strategies and their cumulative effect could be. The circuit court affirmed a \$1.27 billion award against Tucker after he lost on summary judgment.

*AMG Capital Management* isn't the only recent appellate court opinion in which the courts have regarded dark pattern techniques as deceptive trade practices. In *Federal Trade Commission v. LeadClick Media*, the Second Circuit confronted "disguised ad" behavior and false testimonials.<sup>94</sup> LeadClick was an internet advertising company, and its key customer was LeanSpa, an internet retailer that sold weight-loss and colon-cleanse products.<sup>95</sup> LeadClick's strategy was to place much of its advertising on web sites that hosted fake news. Many of the advertisements it placed purported to be online news articles but they were in fact ads for LeanSpa's products. The supposed articles included photos and bylines of the phony journalists who had produced the stories extolling the virtues of LeanSpa's products. As the court explained, these "articles generally represented that a reporter had performed independent tests that demonstrated the efficacy of the weight loss products. The websites also frequently included a 'consumer comment' section where purported 'consumers' praised the products. But there were no consumers commenting – this content was invented."<sup>96</sup> The Second Circuit thought it was self-evident that these techniques were unlawfully deceptive, reaching that conclusion after articulating the applicable legal standard.<sup>97</sup> Again, the court lacked the vocabulary of dark patterns, and also lacked data about their efficacy, but it still regarded the issue as straightforward. The Second Circuit's decision echoed a First Circuit decision from the same year, *Fanning v. Federal Trade Commission*, in which that court treated a defendant's incorrect implication that content was user-generated as a deceptive practice in trade.<sup>98</sup>

A recent deceptive conduct F.T.C. action against Office Depot is instructive as to the agency's current thinking. In that complaint, the F.T.C. alleged that Office Depot and its corporate partner, Support.com, were falsely informing consumers that their computers were

---

<sup>93</sup> Id. at 422 (noting the densely packed text); 423 (noting that consumers had to take affirmative action to avoid the renewal option and that there would be subsequent renewals after that); 423-424 ("nothing in the fine print explicitly states that the loan's 'renewal' would be the automatic consequence of inaction. Instead, it misleadingly says that such renewal must be 'accepted,' which seems to require the borrower to perform some affirmative action."); 424 (noting that "between the sentence that introduces the decline-to-renew option and the sentences that explain the costly consequences of renewal, there is a long and irrelevant sentence about what happens if a pay date falls on a weekend or holiday").

<sup>94</sup> F.T.C. v. LeadClick Media, LLC, 838 F.3d 158 (2d. Cir. 2016).

<sup>95</sup> Id. at 163.

<sup>96</sup> Id at 163-64.

<sup>97</sup> Id. at 168. Though it is not relevant to the portions of the opinion cited here, there is some unfortunate sloppiness in the *LeadClick* opinion. In a couple of instances, the opinion conflates unfair and deceptive practices in trade. See, e.g., *id.* at 168 (erroneously stating that a deceptive practices suit must show that the injury to consumers is not reasonably avoidable by the consumers and is not outweighed by countervailing benefits to consumers or to competition, with a statutory citation that explicitly references the law regarding unfair practices, not deceptive practices). The law is clear that these are not elements of deceptive practices claims. See *Cyberspace.Com*, 453 F.3d at 1199 n.2.

<sup>98</sup> *Fanning*, 821 F.3d at 171-73.



infected with malware and then selling them various fixes for non-existent problems.<sup>99</sup> Office Depot and Support.com were apparently employing misleading software that convinced consumers to pay money for virus and malware removal services they did not need.

Advertisements and in-store sales associates encouraged customers to bring their computers to Office Depot for free “PC Health Checks.” When a customer did so, Office Depot employees would ask consumers whether they had any of the following four problems with their computer: (1) frequent pop-up ads, (2) a computer that was running slowly; (3) warnings about virus infections, or (4) a computer that crashed frequently.<sup>100</sup> If the answer to any of those questions was yes, the employees were to check a corresponding box on the first screen of the Health Check software. The computers then had their systems scanned by Office Depot employees using the Support.com software. Customers were led to believe that the process of scanning the computers was what generated subsequent recommendations from Office Depot employees about necessary fixes, such as virus and malware removal services. In fact, the scanning process was irrelevant for the purposes of generating such recommendations. The only relevant factors for generating recommendations were the responses to the first four questions that the employee asked the customer.<sup>101</sup>

Office Depot strongly encouraged its affiliated stores to push customers towards the PC Health Checks and allegedly expected a high percentage (upwards of 50%) of these Health Checks to result in subsequent computer repairs. Various store employees raised internal alarms about the software, noting that it was flagging as compromised computers that were working properly. These internal complaints evidently were ignored at the C-suite level. Eventually a whistle-blower called reporters at a local Seattle television station. The station had its investigative reporters purchase brand new computers straight from the manufacturers and then bring those computers into Office Depot for PC Health Checks. In several cases, the Support.com software indicated that virus and malware removal was needed. Oops. The journalists’ revelation resulted in an F.T.C. investigation and Office Depot quickly pulled the plug on its PC Health Check software. The companies settled with the F.T.C., agreeing to pay \$25 million (in the case of Office Depot) and \$10 million (in the case of Support.com) to make the case go away, albeit with no admission of wrongdoing on the part of either company.<sup>102</sup>

Several aspects of the deception in *Office Depot* resemble dark patterns. The entire computer scanning process was an example of aesthetic manipulation / hidden information designed to make the customer think that something other than their answers to the first four questions (yes, I see annoying pop-up ads) were driving the company’s recommendations about

---

<sup>99</sup> Federal Trade Commission v. Office Depot, Complaint for Permanent Injunction and Other Equitable Relief, Case No. 9-19-cv-80431 (S.D. Fla. Mar. 27, 2019), available at [https://www.FTC.gov/system/files/documents/cases/office\\_depot\\_complaint\\_3-27-19.pdf](https://www.FTC.gov/system/files/documents/cases/office_depot_complaint_3-27-19.pdf).

<sup>100</sup> *Id.* at 10.

<sup>101</sup> Note the similarity between Office Depot’s computer scans and our bogus calculation of each subject’s “privacy propensity score” in the experiment.

<sup>102</sup> Federal Trade Commission v. Office Depot, Stipulated Order for Permanent Injunction and Monetary Judgment, Case No. 9-19-cv-80431-RLR (S.D. Fla. Mar. 28, 2019), available at [https://www.FTC.gov/system/files/documents/cases/office\\_depot\\_stipulated\\_order\\_3-29-19.pdf](https://www.FTC.gov/system/files/documents/cases/office_depot_stipulated_order_3-29-19.pdf) (Office Depot settlement) and [https://www.FTC.gov/system/files/documents/cases/office\\_depot\\_-\\_support.com\\_stipulated\\_order\\_3-29-19.pdf](https://www.FTC.gov/system/files/documents/cases/office_depot_-_support.com_stipulated_order_3-29-19.pdf) (Support.com settlement); Michelle Singletary, *Office Depot and Support.com to Pay \$35 Million to Settle Charges of Tech Support Scam*, WASH. POST, Mar. 28, 2019.

necessary repairs. There is also a clear bait-and-switch component to the allegations against Office Depot – customers thought they were getting a helpful and free diagnostic from a respected retailer. Instead, they were opening themselves up to a deceitful way for Office Depot to upsell services that many customers did not need. This was done via a mediated online interface employed in brick-and-mortar retail outlets.

Critically, in deciding what constitutes a deceptive practice in trade, the fact that many consumers wind up with terms, goods, or services they do not want strongly suggests that the seller has engaged in deception. That is a key take-away from another Ninth Circuit case, *Cyberspace.com*.<sup>103</sup> In that case, a company mailed personal checks to potential customers, and the fine print on the back of those checks indicated that by cashing the check the consumers were signing up for a monthly subscription that would entitle them to internet access. Hundreds of thousands of consumers and small businesses cashed the checks, but less than one percent of them ever utilized the defendant's internet access service.<sup>104</sup> That so many consumers had been stuck with something they didn't desire and were not using was "highly probative," indicating that most consumers "did not realize they had contracted for internet service when they cashed or deposited the solicitation check."<sup>105</sup> Courts considering F.T.C. section 5 unfairness suits, discussed below, embrace the same kind of evidence and reasoning.<sup>106</sup> By the same logic, if it appears that a large number of consumers are being dark patterned into a service they do not want (as occurred in our experiment) then this evidence strongly supports a conclusion that the tactics used to produce this assent are deceptive practices in trade.

There is less clear case law surrounding the F.T.C.'s use of section 5 from which to construct a profile of what conduct is "unfair." In the overwhelming majority of enforcement actions, companies choose to settle with the Commission, entering into binding settlement agreements, rather than challenge the commission in court or administrative proceedings.<sup>107</sup> In the absence of judicial decisions; however, consent decrees and other F.T.C. publications have guided companies in interpreting the expected standards of behavior and ensuring their continued compliance with the law.<sup>108</sup>

In 1980, the F.T.C. laid out the test that is still currently utilized to find an act or practice "unfair." Under this test, an unfair trade practice is one that 1) causes or is likely to cause substantial injury to consumers 2) is not reasonably avoidable by consumers themselves and 3) is not outweighed by countervailing benefits to consumers or competition.<sup>109</sup> This three-part test is now codified in section 5(n) of the F.T.C. Act.

Generally, the "substantial injury" prong focuses on whether consumers have suffered a pecuniary loss. Monetary harm can come from the coercion of consumers into purchasing

---

<sup>103</sup> 453 F.3d at 1196.

<sup>104</sup> *Id.* at 1199.

<sup>105</sup> *Id.* at 1201.

<sup>106</sup> See, e.g., *F.T.C. v. Direct Benefits Group, LLC*, 2013 WL 3771322, Case No. 6:11-cv-1186-Orl-28TBS, at \*14 (M.D. Fla. July 18, 2013).

<sup>107</sup> Daniel J. Solove & Woodrow Hartzog, *The F.T.C. and the New Common Law of Privacy*, 114 *COL. L. REV.* 583 (2014).

<sup>108</sup> *Id.*

<sup>109</sup> F.T.C. Policy Statement on Unfairness, 104 F.T.C. 949 (1984).

unwanted goods, or other incidental injuries that come as a result of the unfair action such as financial harm from identity theft. Notably, a harm's substantiality can derive from its collective effect on consumers, as the F.T.C. notes "an injury may be sufficiently substantial, however, if it does a small harm to a large number of people."<sup>110</sup>

The next prong of the three-part unfairness test is that the injury must not be one that the consumer could have reasonably avoided. This prong is grounded in the belief that the market will be self-correcting and that consumers will learn to avoid companies that utilize unfair practices. Those practices that "prevent consumers from effectively making their own decisions," run afoul of this prong, even if they merely hinder free market decisions, and fall short of depriving a consumer of free choice. For reasonable consumers to avoid harm, particularly in the case of a nonobvious danger, they must also be aware of the possible risk.

The cost-benefit analysis prong of the unfairness test ensures that companies are only punished for behaviors that produce "injurious net effects." There are, as the Commission notes, inevitable trade-offs in business practices between costs and benefits for consumers, and as such certain costs may be imposed on consumers, provided they are balanced by legitimate benefits. Broader societal burdens are also accounted for in this equation, as are the potential costs that a remedy would entail. Additionally, the Commission looks to public policy considerations as part of this analysis to help establish the existence and weight of injuries and benefits that are not easily quantified.

A few cases that resemble dark pattern conduct were brought on unfairness grounds as well as deception. A number of these F.T.C. cases involve unsavory billing practices. One example is *F.T.C. v. Bunzai Media Group, Inc.*, a case in which the F.T.C. secured a settlement of upwards of \$73 million after alleging both deceptive and unfair practices.<sup>111</sup> In that case the F.T.C. asserted that the defendants' skin-care companies were using a host of dark patterns, including deceptive pop-up ads that stopped consumers from navigating away from a web site without accepting an offer, small print at the very end of a transaction that were in tension with marketing claims used in larger, bold print, and pricing plans that quickly converted "risk-free trials" into renewing monthly subscriptions and were onerous to cancel.<sup>112</sup> The F.T.C.'s more recent suit against Triangle Media involved some similar sales tactics, plus a nasty surprise – at the end of the transaction to set up the "free trial," the defendants used misleading web site text to create the false impression that the transaction was not complete until customers signed up for a second free trial for an entirely different product, and they would be signed up for costly monthly subscriptions to both by clicking on the "complete checkout" button.<sup>113</sup> This case too was brought under both prongs of section 5 – deception and unfairness.

---

<sup>110</sup> *Id.*

<sup>111</sup> 2016 WL 3922625, at \*5, Case No. CV 15-4527-GW(PLAx) (C.D. Cal. July 19, 2016).

<sup>112</sup> *F.T.C. v. Bunzai Media Group*, Case No. CV 15-4527-GW(PLAx), First Amended Complaint for Permanent Injunction and Other Equitable Relief, (C.D. Cal. Oct. 9, 2015), available at <https://www.FTC.gov/system/files/documents/cases/151009bunzaicmpt.pdf>

<sup>113</sup> *F.T.C. v. Triangle Media Corp.*, 2018 WL 4051701, Case No. 18cv1388-MMA (NLS) (S.D. Cal. Aug. 24, 2018).

*F.T.C. v FrostWire, LLC*,<sup>114</sup> is another case involving alleged unfairness as well as deception, this time with respect to the default settings of a peer-to-peer file sharing service that caused users to share more media than they were lead to believe. The F.T.C. pointed to the obstructionist defaults of the program, which made it exceptionally burdensome for a consumer to prevent all of her files from being shared. As described in the complaint "a consumer with 200 photos on her mobile device who installed the application with the intent of sharing only ten of those photos first had to designate all 200 ... as shared, and then affirmatively unshare each of the 190 photos that she wished to keep private." This user interface presents a classic roach motel employing preselection.

These cases notwithstanding, there is little case law discussing unfairness and dark patterns in depth, especially in comparison to the development of the deceptive acts and practices precedents. Worse still, the leading appellate unfairness case is a Ninth Circuit unpublished disposition that lacks precedential value. The court concluded in that case, for example, that it was unfair conduct for material language to appear in blue font against a blue background on an "otherwise busy" web page.<sup>115</sup>

Many of the dark patterns discussed earlier could be characterized in a manner to frame the injury as a consumer entering into a transaction they otherwise would have avoided, therefore falling squarely into the current conception of substantial injury. That said, there may be hurdles in conceptualizing dark patterns in a way that fulfills the "unavoidability" prong. When the use of dark patterns is extreme, capitalizing on consumer cognitive bias to the extent that it can be shown to overwhelm their ability to make a free decision, there should be no problem satisfying this prong. At first blush, the milder the use of dark patterns, the more difficult it will be to characterize the harm as unavoidable, particularly when not applied to any exceptionally vulnerable subsets of consumers. On the other hand, our data suggests that milder dark patterns are – if anything – harder to avoid, because of their potent combination of subtlety and persuasive ability.

To summarize, there is an emerging body of precedent in which the federal courts have viewed the F.T.C. as well within its rights to pursue companies that deploy dark patterns online. Among the techniques identified in the taxonomy, false testimonials, roach motels, hidden costs, forced continuity, aesthetic manipulation, preselection, trick questions, and disguised ads have already formed the basis for violations of the F.T.C.'s prohibition on deceptive acts in trade. Techniques that also employ deception, such as false activity messages, sneaking into the basket, bait and switch, forced registration, and scarcity techniques would seem to fall straightforwardly within the parameters of the existing law. Other techniques, like nagging,

---

<sup>114</sup> Complaint for Permanent Injunction and Other Equitable Relief, Oct. 7, 2011, available at 2011 WL 9282853.

<sup>115</sup> *F.T.C. v. Commerce Planet, Inc.*, 642 Fed.Appx. 680, 682 (9th Cir. Mar. 3, 2016). The district court's opinion, which is published, and which was affirmed in this respect by the Ninth Circuit, provides more detail. See *F.T.C. v. Commerce Planet, Inc.*, 878 F. Supp.2d 1048, 1066 (C.D. Cal. 2012) ("As placed, the disclosure regarding OnlineSupplier's negative option plan is difficult to read because it is printed in the smallest text size on the page and in blue font against a slightly lighter blue background at the very end of the disclosure. The disclosure is also not placed in close proximity to the 'Ship My Kit!' button and placed below the fold. It is highly probable that a reasonable consumer using this billing page would not scroll to the bottom and would simply consummate the transaction by clicking the 'Ship My Kit!' button, as the consumer is urged to do by the message at the top left: 'You are ONE CLICK AWAY from receiving the most up-to-date information for making money on eBay!'").

price comparison prevention, intermediate currency, toying with emotion, or confirmshaming would probably need to be challenged under section 5's unfairness prong. We were not able to find cases that shed light on whether nagging, toying with emotion, and confirmshaming are lawful. In any event, this survey of the existing precedents suggests that the law restricting dark patterns does not need to be invented; to a substantial degree it's already present.

State unfair competition laws largely track their federal counterpart. There has been far less enforcement activity under these laws targeting dark patterns than there has been under the applicable federal regime. As a result, the law is underdeveloped, and few state cases have broken new ground. An exception is *Kulsea v. PC Cleaner, Inc.*,<sup>116</sup> a case brought under California's unfair competition law that predated, and in many ways anticipated, the F.T.C.'s suit against Office Depot. The allegations against PC Cleaner were that the firm's software indicated that there were harmful bugs on the machine that could be addressed via the purchase of the full version of the software.

Another instructive state law case is *In re Lenovo Adware Litigation*.<sup>117</sup> That class action case is a sort of split-decision where dark patterns are concerned. Lenovo pre-installed adware on computers that it sold to customers, hiding the software deep within the computers' operating system so it would be difficult to detect and remove. Consumers were given just one chance to remove the software the first time they opened their internet browser, and retaining the software was the default option. Lenovo thus employed preselection, alongside arguable bait-and-switch and hidden costs. A claim brought under New York's consumer protection law, which prohibits deceptive trade practices, was dismissed because the plaintiffs failed to show that they suffered an actual injury, such as a pecuniary harm.<sup>118</sup> In the court's view, this lack of pecuniary harm did not justify dismissing the plaintiffs' claims under California state unfair competition law, given that the adware negatively affected the performance of the laptops, and that the installation of the adware was peculiarly within Lenovo's knowledge, material, and a fact that went undisclosed to consumers.<sup>119</sup> The case ultimately settled for more than \$8 million.<sup>120</sup>

## B. Other Relevant Federal Frameworks

Some enforcement efforts that target dark patterns could be done through the Consumer Financial Protection Bureau (C.F.P.B.), which has the authority to regulate "abusive conduct," at least within the banking and financial services sector. The C.F.P.B. abusive conduct definition is arguably more expansive than the unfair conduct that can be regulated by the F.T.C. An abusive practice, per 12 U.S.C. § 5531 is one that:

---

<sup>116</sup> 2014 WL 12581769, NO. CV 12-0725 FMO (ANx), (C.D. Cal. Feb. 10, 2014).

<sup>117</sup> 2016 WL 6277245 (N.D. Cal. Oct. 27, 2016).

<sup>118</sup> *Id.* at \*10.

<sup>119</sup> *Id.* at \*11-\*14.

<sup>120</sup> *In re Lenovo Adware Litigation*, 2019 WL 1791420, at \*6 Case No. 15-md-02624-HSG (N.D. Cal. Apr. 24, 2019).

(1) materially interferes with the ability of a consumer to understand a term or condition of a consumer financial product or service; or

(2) takes unreasonable advantage of -

(A) a lack of understanding on the part of the consumer of the material risks, costs, or conditions of the product or service;

(B) the inability of the consumer to protect the interests of the consumer in selecting or using a consumer financial product or service; or

(C) the reasonable reliance by the consumer on a covered person to act in the interests of the consumer.

This provision would seemingly cover the exploitation of the cognitive biases of consumers in order to manipulate them into making a decision that may not be in their best interests.

Another relevant federal law is the Restore Online Shoppers' Confidence Act (ROSCA).<sup>121</sup> ROSCA makes it unlawful for a third party seller to charge customers absent a clear and conspicuous disclosure of the transaction's material terms, informed consent, and an affirmative step by the consumer indicating willingness to enter into the transaction with the third party.<sup>122</sup> This law was aimed at the problem of consumers unwittingly being signed up for a subscription to a third party's good or service immediately after entering into a desired transaction with a vendor, where the third party would use the payment information that the consumer had already inputted. The F.T.C. enforces ROSCA in a manner similar to its section 5 enforcement, and ROSCA squarely addresses certain types of bait-and-switch dark patterns, which often employed hidden costs and forced continuity schemes.

### C. Contracts and Consent

In his 2018 book, Woodrow Hartzog advanced the argument that contractual consent secured via pernicious forms of dark patterns or other deceptive designs should be deemed invalid as a matter of law.<sup>123</sup> Hartzog's argument is built on a series of powerful anecdotes, and the eye-opening data we present here buttresses his bottom line. In our view, Hartzog has it mostly right. The hard part, however, is determining how to tell whether a dark pattern is egregious enough to disregard a consumer's clicking of an "I agree" button. Hartzog's book spends just a few pages developing that particular argument, so there is more theoretical and doctrinal work to be done.

The law's deference to contractual arrangements is premised on a belief that private ordering that commands the mutual assent of the parties makes them better off than the alternative of mandatory rules whose terms are set by the government. The more confidence we have that a contractual arrangement is misunderstood by one of the parties and does not serve the expressed interests of that party, the less reason there is to let the terms of a relationship be set by contract law. To put matters in terms of an influential argument recently

---

<sup>121</sup> 15 U.S.C. §§8401-05.

<sup>122</sup> 15 U.S.C. § 8402.

<sup>123</sup> HARTZOG, *supra* note 35, at 212-13.

advanced by Rob Kar and Peggy Radin, assent procured mostly via the use of dark patterns doesn't form contracts; it forms pseudo-contracts.<sup>124</sup> Those shouldn't bind the signatories.

At first blush, hostility to consent induced by dark patterns does not appear to be the direction that the contracts case law has been going of late, though a large part of the problem may be the absence of evidence like the data that our study reveals. *Williams v. Affinion Group, LLC*,<sup>125</sup> is a key recent case. In *Williams* a confusing user interface was employed by the defendant, Trilegiant, to sign up consumers for membership club purchases while consumers were in the process of shopping for goods and services on sites like Priceline.com.<sup>126</sup> The consumers were given a discount on their Priceline purchase if they signed up for a membership in one of the defendant's clubs, and if they did so they would be billed \$10 to \$20 monthly for said membership until the consumer cancelled it.<sup>127</sup> As the Second Circuit described it:

To snare members, Trilegiant allegedly designs its enrollment screens to appear as confirmation pages for the legitimate, just-completed transaction, so that the customer is unaware of having registered to buy and new and completely different product. Trilegiant's cancellation and billing process allegedly prolongs the fraud. To cancel a subscription, the customer must first discover the monthly billing on a credit card statement and call Trilegiant's customer service; Trilegiant's representatives then attempt to keep members enrolled as long as possible, either through promotion of the program's benefits or delay in the cancellation process.<sup>128</sup>

To be clear, not everything described above is a dark pattern, but some of those steps – the disguised ad, the roach motel, the forced continuity, and the nagging – would qualify. The district court's opinion helpfully reproduced the text of Trilegiant's user interface, albeit with much of the text too small to read.<sup>129</sup> From that text and the lower court opinion it appears the plaintiffs were arguing that the deceptive conduct was evident from a glance at the screenshots.

To the *Williams* court, there was insufficient evidence that this conduct vitiated consent. The plaintiffs produced an expert witness, a marketing scholar, who testified that the user interface "was designed to result in purchases of Trilegiant's services without awareness of those purchases,"<sup>130</sup> and that the disclosures were designed "so that they would not be seen or understood."<sup>131</sup> The plaintiff's also argued that the relevant terms of the program were buried in "miniscule fine print."<sup>132</sup>

---

<sup>124</sup> Robin Bradley Kar & Margaret Jane Radin, *Pseudo-Contract and Shared Meaning Analysis*, 132 HARV. L. REV. 1135, 1192-1201 (2019).

<sup>125</sup> 889 F.3d 116 (2d. Cir. 2018).

<sup>126</sup> *Id.* at 117.

<sup>127</sup> *Id.* at 120.

<sup>128</sup> *Id.*

<sup>129</sup> *In re Trilegiant Corp.*, 2016 WL 8114194, at \*2 (D. Conn. Aug. 23, 2016).

<sup>130</sup> *Williams*, 889 F.3d at 123.

<sup>131</sup> *Id.* at 122.

<sup>132</sup> *Id.* at 122.

The plaintiff made two key mistakes that, from the Second Circuit’s perspective, warranted the district court’s decision to grant the defendant’s summary judgment motion. First, the expert witness does not appear to have presented any data about consumer confusion – his statements about the interface design and Trilegiant’s likely intentions were conclusory and not supported by evidence in the record.<sup>133</sup> Second, the plaintiffs did not argue that the plaintiffs were confused as a result of ambiguous language or design.<sup>134</sup> In short, the *Williams* opinion leaves the door ajar for class action suits against ecommerce firms that employ dark patterns, provided the proof of consumers being confused or tricked into paying for goods and services they do not want employs the kind of rigorous randomization-based testing that we present here.

The contract doctrine of undue influence provides the most promising existing framework for efforts to curtail dark patterns. Under the Restatement (Second) of Contracts, “undue influence is unfair persuasion of a party who is under the domination of the person exercising the persuasion or who by virtue of the relation between them is justified in assuming that that person will not act in a manner inconsistent with his welfare.”<sup>135</sup> Comment b of the Restatement emphasizes further that the “law of undue influence ... affords protection in situations where the rules on duress and misrepresentation give no relief. The degree of persuasion that is unfair depends on a variety of circumstances. The ultimate question is whether the result was produced by means that seriously impaired the free and competent exercise of judgment. Such factors as the unfairness of the resulting bargain, the unavailability of independent advice, and the susceptibility of the person persuaded are circumstances to be taken into account in determining whether there was unfair persuasion, but they are not in themselves controlling.”<sup>136</sup> Undue influence renders a contract voidable by the influenced party.<sup>137</sup>

Applying this rubric, it should not be controversial to assert that packages of dark patterns like the ones employed in our experiment seriously impaired the free and competent exercise of judgment. That seems to be their purpose and effect, as our data show. The harder doctrinal question is whether a consumer and the typical firm that employs dark patterns establishes satisfies either the domination or relationship part of the Restatement test.

The case law suggests that some courts construe the relationship language broadly. In one prominent case, a chiropractor convinced his patient to sign a form indicating that the

---

<sup>133</sup> *In re Trilegiant Corp.*, 2016 WL 8114194, at \*11 n.3.

<sup>134</sup> *Williams*, 889 F.3d at 123 (“[T]o show that customers may have been misled, the plaintiff must produce evidence that particular statements are deceptive when considered in context. These plaintiffs have not attempted to do so. This is not a case involving confusing text; instead, the plaintiffs’ primary contention is that the appearance of an enrollment offer in the course of a separate e-merchant transaction was itself inherently deceptive because it led customers to believe that Trilegiant’s products were associated with or offered by the e-merchant. . . . [T]he plaintiffs’ theory that misleading enrollment pages deceived them into believing they were enrolling in something other than a discount club membership is entirely inconsistent with the record evidence that individual plaintiffs were unaware they enrolled in anything to begin with.”).

<sup>135</sup> RESTATEMENT (SECOND) OF CONTRACTS § 177 (1981).

<sup>136</sup> *Id.* at § 177 comment b.

<sup>137</sup> *Rich v. Fuller*, 666 A.2d 71, 76 (Maine 1995).



patient would pay for the services in full even if her insurance company elected not to cover them.<sup>138</sup> When the patient objected, saying that she could not afford to pay out of pocket, the chiropractor told her “that if her insurance company said they would take care of her, they would. He told her not to worry.”<sup>139</sup> These statements induced the patient to sign. The court granted summary judgment to the chiropractor against the patient’s undue influence claim, and the appellate court reversed. From the appellate court’s perspective, these statements uttered in the context of this medical treatment relationship was enough for a reasonable jury to conclude that undue influence had occurred.<sup>140</sup> The majority brushed aside the concerns of a dissenting judge, who accused the majority of invalidating a contract over “nothing more than the urging, encouragement, or persuasion that will occur routinely in everyday business transactions.”<sup>141</sup> Another leading case where the court similarly reversed a summary judgment motion involved a relationship between a widow and her long-time friend who was also an attorney.<sup>142</sup>

In influential publications, Jack Balkin and Jonathan Zittrain have proposed that digital platforms like Facebook, Google, Microsoft, and Amazon should owe fiduciary duties to their customers.<sup>143</sup> If such a proposal were implemented, then the use of effective dark patterns by these platforms would render any consent procured thereby voidable by the customer. This result follows because the law generally presumes undue influence in those instances where a fiduciary owes a duty to a client and the fiduciary benefits from a transaction with its client.<sup>144</sup>

Even without embracing Balkin and Zittrain’s information fiduciary theory,<sup>145</sup> dark patterns could be voidable under the domination theory referenced in the Restatement. There is some fuzziness around the precise meaning of domination in the case law. Circumstantial evidence is plainly adequate to prove undue influence.<sup>146</sup> A classic undue influence case describes domination as a kind of “overpersuasion” that applies pressure that “works on mental, moral, or emotional weakness to such an extent that it approaches the boundaries of coercion.”<sup>147</sup> As the court emphasized, “a confidential or authoritative relationship between the

---

<sup>138</sup> Gerimonte v. Case, 712 P.2d 876 (Wash. App. 1986).

<sup>139</sup> Id. at 877.

<sup>140</sup> Id. at 879.

<sup>141</sup> Id. at 880 (Scholfield, C.J., dissenting).

<sup>142</sup> Goldman v. Bequai, 19 F.3d 666, 669 (D.C. Cir. 1994).

<sup>143</sup> See Jack M. Balkin, *Information Fiduciaries and the First Amendment*, 49 U.C. DAVIS L. REV. 1183 (2016); Jack M. Balkin & Jonathan Zittrain, *A Grand Bargain to Make Tech Companies Trustworthy*, THE ATLANTIC, Oct. 3, 2016, available at <https://www.theatlantic.com/technology/archive/2016/10/information-fiduciary/502346/>.

<sup>144</sup> See, e.g., Matlock v. Simpson, 902 S.W.2d 385, 386 (Tenn. 1995). In those situations the fiduciary must demonstrate the substantive fairness of the underlying transaction to defeat a claim of undue influence.

<sup>145</sup> For a critique of Balkin & Zittrain’s proposal see Lina Khan & David Pozen, *A Skeptical View of Information Fiduciaries*, 133 HARV. L. REV. \_\_\_\_ (forthcoming 2019), available at [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3341661](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3341661).

<sup>146</sup> Nichols v. Estate of Tyler, 910 N.E.2d 221 (Ind. App. 2009); In re Cheryl E., 207 Cal. Rptr. 728, 737 (Cal. App. 1984).

<sup>147</sup> Odorizzi v. Bloomfield Sch. Dist., 54 Cal. Rptr. 533, 539 (Cal. App. 1966).

parties need not be present when the undue influence involves unfair advantage taken of another's weakness or distress."<sup>148</sup> In the court's judgment, undue influence could arise when "a person of subnormal capacities has been subjected to ordinary force or a person of normal capacities subjected to extraordinary force."<sup>149</sup> None of the cases suggest that domination requires a certainty that the dominated party will do the dominant party's bidding.

Nearly quadrupling the percentage of consumers who surrender and agree to waive their rights through non-persuasive tactics like nagging, confusion, hidden costs, or roach motels could satisfy the domination test, particularly when those tactics are unleashed against relatively unsophisticated users. Indeed, in trying to determine whether a tactic amounts to undue influence, courts have emphasized factors such as "limited education and business experience"<sup>150</sup> as well as the uneven nature of the exchange in terms of what the party exercising influence gave and received.<sup>151</sup> Similarly, the Restatement identifies "the unfairness of the resulting bargain, the unavailability of independent advice, and the susceptibility of the person persuaded" as the relevant considerations.<sup>152</sup> Treating highly effective dark patterns as instances of domination-induced undue influence would amount to an extension of the doctrine, but it's an extension consistent with the purpose of the doctrine. Furthermore, the availability of quantifiable evidence about the effects of particular dark patterns addresses lingering problems of proof that might otherwise make judges skeptical of the doctrine's application. In short, there are sensible reasons to think that the use of dark patterns to secure a consumer's consent can render that consent voidable by virtue of undue influence.

To push the argument further, there are a number of instances in which the existence of consent is necessary in order for the sophisticated party to a transaction to engage in conduct that would otherwise be unlawful. We identify three such statutory frameworks here. The first is electronic communications law. It is unlawful to intercept an electronic communication (such as a phone call or an email) without the consent of the parties to a communication.<sup>153</sup> Failure to secure consent has given rise to civil suits under this provision of the Electronic Communications Privacy Act and its state law equivalents.<sup>154</sup> There is a strong argument to be made that consent secured via dark patterns is not adequate consent under these statutes, thereby opening up parties that intercept such communications to substantial liability, especially in cases where

---

<sup>148</sup> *Id.* at 540. Some, though not all, of the factors relevant to identifying overpersuasion are common in certain forms of dark patterns, such as "discussion of a transaction at an unusual or inappropriate time," "insistent demand that the business be finished at once," "the use of multiple persuaders by the dominant side against a single servient party," and "the absence of third-party advisers to the servient party." As the court explained, "[i]f a number of these elements are simultaneously present, the persuasion may be characterized as excessive." *Id.* at 541.

<sup>149</sup> *Id.* at 541.

<sup>150</sup> *Delaney v. Delaney*, 402 N.W.2d 701, 705 (S.D. 1987).

<sup>151</sup> *Goldman*, 19 F.3d at 675.

<sup>152</sup> RESTATEMENT (SECOND) OF CONTRACTS § 177 comment b.

<sup>153</sup> 18 U.S.C. § 2511(2)(d).

<sup>154</sup> See, e.g., *Deal v. Spears*; *In re Yahoo Mail Litigation*, 7 F. Supp.3d 1016 (N.D. Cal. 2014); *In re Google Inc. Gmail Litigation*, 2013 WL 5423918 (N.D. Cal. Sep. 26, 2013) (No. 13-MD-02430-LHK).

large numbers of communications have been intercepted, such as controversies involving automated content analysis of emails.

Illinois' unique Biometric Identification Privacy Act (BIPA) places similar emphasis on the consent requirement. It requires firms that process the biometric information of consumers to obtain their explicit consent before doing so. The Illinois law sets a high threshold for what counts as adequate consent – firms must inform customers of the fact that biometric information is being collected and stored, the reason for collection, use, and storage, and the duration of storage.<sup>155</sup> The law has produced an avalanche of class action litigation, directed at firms that analyze fingerprints, facial geometry in photos, voiceprints, or other biometric information. In the first half of 2019 new class action suits under BIPA were being filed at a rate of approximately one per day.<sup>156</sup> This rate of new class actions is driven in part by the availability of minimum statutory damages under the statute and the determination by the Illinois Supreme Court that it is not necessary to demonstrate an actual injury in order to have standing to sue under the statute in state court.<sup>157</sup> As ecommerce firms increasingly recognize the scope of their potential exposure to BIPA damages, many have done more to provide the disclosure boxes required by the statute. To the extent that they do so via a disclosure or consent-extracting mechanism that employs dark patterns, the courts could well deem those interfaces (and the “consent” produced thereby) inadequate as a matter of law, opening up the firms that employ those mechanisms subject to very significant liability.<sup>158</sup>

A relevant, but not heavily utilized, law exists in California as well. That state enacted a law in 2009 that can be used to aim squarely at forced continuity dark patterns. The law would “end the practice of ongoing charging of consumer credit or debit cards or third party payment accounts without the consumers’ explicit consent for ongoing shipments of a product or ongoing deliveries of service.”<sup>159</sup> Recall Sony’s use of a roach motel to substantially thwart the wishes of PlayStation Plus users who wish to avoid a renewing subscription. There is a very plausible argument that Sony’s obstruction scheme, and ones like it, fall short of the explicit consumer consent standard required by California law. Without stretching the meaning of the statute’s words it is easy to imagine significant class action exposure for Sony.

#### D. Line Drawing

We expect that most readers will have some sympathy for the idea that dark patterns could be so pervasive in a particular context as to obviate consent. But the hard question, and one readers have probably had on their minds as they read through the preceding pages, is “where does one draw the line?” We would readily concede that some dark patterns are too minor to warrant dramatic remedies like contractual rescission, and some do not warrant a

---

<sup>155</sup> 740 Ill. Comp. Stat. 14/20(2).

<sup>156</sup> See Seyferth Shaw LLP, *Biometric Privacy Class Actions by the Numbers: Analyzing Illinois’ Hottest Class Action Trend*, July 2, 2019, available at <https://www.jdsupra.com/legalnews/biometric-privacy-class-actions-by-the-48938/>.

<sup>157</sup> *Rosenbach v. Six Flags Entertainment Corp.*, \_\_\_ N.E.3d \_\_\_, available at 2019 WL 323902 (Ill. Jan. 25, 2019).

<sup>158</sup> For a discussion of liability under these provisions of federal and state wiretap acts and BIPA, see Lior Jacob Strahilevitz & Matthew B. Kugler, *Is Privacy Policy Language Irrelevant to Consumers?*, 45 J. LEGAL STUD. S69 (2016).

<sup>159</sup> Cal. Bus. & Prof. Code § 17600.

regulatory response of any sort. Small dosages of nagging, intermediate currency, toying with emotion, and confirmshaming may be close to benign and could even be mildly beneficial in limited contexts.<sup>160</sup> Policing them aggressively is unlikely to be cost-justified. At the same time, an “I know it when I see it” approach to dark patterns creates uncertainty, notice problems, and raises the specter of unequal enforcement.

We believe there is a better way forward. In our view, a quantitative approach to identifying dark patterns could be workable and offers many of the benefits of bright-line rules in general. More precisely, where the kind of A/B testing that we discuss above reveals that a particular interface design or option set more than doubles the percentage of users who wind up “consenting” to engage in a consumer transaction, the company practice at issue could be deemed presumptively an unfair or deceptive practice in trade. In the scenarios tested in our experiment, both the mild dark patterns and the hard dark patterns made it more likely than not that consumers were electing not to decline a service on the basis of the choice architecture employed rather than on the basis of innate demand for the service at issue. The “more likely than not” standard is widely employed in civil litigation over torts and other kinds of liability, and it could work well in this context too, ideally with the F.T.C. and academics working hand in hand to replicate high-quality research that quantifies the effects of particular manipulations. As a statistical matter, each individual research subject in our study who was signed up for the data protection plan was more likely than not to have done so because of the dark pattern rather than because of underlying demand for the service being offered.

Admittedly, one challenge here is to develop a neutral baseline against which the A/B testing can occur. With respect to straightforward linguistic choices, that sometimes will be easy. It should not be hard to generate consensus around the idea that a simple Yes / No or Accept / Decline prompt is neutral, provided the choices are presented with identical fonts, colors, font sizes, and placement. Things get more challenging when legal decision-makers must determine whether two, three, four, or five options is neutral, and that is an inquiry that is easier to answer with the benefit of data than it is in the abstract. In close cases, dueling experts may testify, and agencies or courts may be called upon to make the same kinds of factual determinations that is the bread and butter of adjudication. Similarly, there may be some challenges in identifying the neutral baseline where aesthetic manipulation is alleged. Here existing practices may help inform rational determinations about how to assess the baseline. Black text on a white background in a common, 12-point font is used widely enough in communication to where a social scientist treating it as a neutral baseline is unlikely to get

---

<sup>160</sup> Take the nagging example. As any parent of verbal kids can attest, a modicum of nagging is entirely tolerable. When kids nag their parents it conveys an intensity of preferences to parents, who may appropriately choose to relent after realizing (on the basis of the persistent nagging) that the requested food, activity, or toy really is very important to the child. That said, the legal system has long recognized that nagging should have its limits. The college student who asks a classmate out on a date once or maybe twice, only to be rebuffed, is behaving within acceptable bounds. As the requests mount in the face of persistent rejection, the questions can become harassment. See U.S. DEPARTMENT OF EDUCATION, OFFICE FOR CIVIL RIGHTS, REVISED SEXUAL HARASSMENT GUIDANCE: HARASSMENT OF STUDENTS BY SCHOOL EMPLOYEES, OTHER STUDENTS, OR THIRD PARTIES 6 (2001) (“[B]ecause students date one another, a request for a date or a gift of flowers, even if unwelcome, would not create a hostile environment. However, there may be circumstances in which repeated, unwelcome requests for dates or similar conduct could create a hostile environment.”). It is plausible that after the fifth or sixth request to turn on notifications is declined, a commercial free speech claim lodged against a policy that prevents further requests becomes weak enough for the restriction to survive *Central Hudson* scrutiny.

skeptical looks. As graphics take priority over text, however, things become more complicated, and adjudicators will need to be on the lookout for efforts by hack social scientists to reach their desired answers by manipulating the supposedly neutral baseline.

Bright line rules are particularly useful in the context of enforcing section 5 of the F.T.C. Act. The Due Process Clause requires that companies be able to anticipate when they will face legal liability and when they will not.<sup>161</sup> Thus, the more clarity exists in section 5, the less likely it becomes that energetic enforcement of the law will conflict with vital constitutional values. Companies are already doing the kind of beta-testing that reveals how effective their interfaces are becoming at changing consumer behavior. To the extent that there is any doubt about a new technique, they can always examine their own design choices and see whether any cross the line.<sup>162</sup>

In short, it would be appropriate for courts to deem instances in which the “more likely than not” test is satisfied as instances in which consumers have not actually consented to the contractual terms at issue and can void the transaction after the fact. To hold otherwise runs the risk of treating consent as a legal fiction, rather than an indication of mutual assent.

In embracing a “more likely than not” rule, we do not mean to rule out the development of multifactor standards that can supplement a rule-based approach. We are not convinced that a “more likely than not” rule is over-inclusive, as long as dark patterns are defined appropriately, but it may be under-inclusive. For example, the “more likely than not” test works very well when the innate preference for a product among consumers stands at 10 or 20%. But when 40 to 50% of consumers would want to sign up for a service or purchase a product, the “more likely than not” test is likely to let too much manipulative conduct survive. In our view, a situation where 40% of consumers opt to buy a service because of innate demand for it and 20% of consumers opt to buy because of a manipulative interface or choice architecture may still be legally problematic. In those settings it will be necessary to develop a standard that supplements the rule we propose.

A multi-factor test for dark patterns that looks to considerations such as (a) evidence of a defendant’s malicious intent or knowledge of detrimental aspects of the user interface’s design, (b) whether vulnerable populations – like less educated consumers, the elderly, or people suffering from chronic medical conditions – are particularly susceptible to the dark pattern, and (c) the magnitude of the costs and benefits produced by the dark pattern would be a good starting point. Evidence about the ex post regret experienced by consumers who found themselves influenced by a dark pattern might be a particularly revealing indicia of the costs. The greater the number of consumers who complained and sought cancellation of a term they didn’t realize they agreed to, or who didn’t utilize a service they found themselves paying for (as the *Cyberspace.com* court indicated),<sup>163</sup> the greater the presumptive magnitude of the associated harm would be. By the same token, if it turned out that consumers were happy ex post with a good or service that a dark pattern manipulated them into obtaining, this would be

---

<sup>161</sup> The leading recent case addressing this issue is *F.T.C. v. Wyndham Worldwide Corp.*, 799 F.3d , 236, 249-59 (3d Cir. 2015).

<sup>162</sup> See Omri Ben-Shahar & Lior Jacob Strahilevitz, *Interpreting Contracts via Surveys and Experiments*, 92 N.Y.U. L. REV. 1753, 1822-24 (2017) (encouraging this kind of beta-testing with consumer contract language).

<sup>163</sup> See *supra* text accompanying notes 103-105.

revealing evidence cutting against liability for the seller. The ends could justify the means for a firm that genuinely was trying to trick consumers for their own good. But here too, (d) experimental evidence about how effective the dark pattern was compared to a neutral choice architecture should be relevant, albeit not dispositive in a multi-factor inquiry. Thus, even the standard we propose would include a sliding scale that is tied to a quantifiable metric.

The “more likely than not” rule also addresses one of the design challenges that legislators seeking to restrict dark patterns have encountered. As we noted at the outset,<sup>164</sup> bipartisan legislation is presently pending in the Senate to prohibit dark patterns. Senate Bill 1084 would treat the activities of any online service with more than 100 million unique users that “design, modify, or manipulate a user interface with the purpose or substantial effect of obscuring, subverting, or impairing user autonomy, decision-making, or choice to obtain consent or user data” as unfair or deceptive practices in trade.<sup>165</sup> At the same time, the legislation recognizes that this open-ended prohibition may leave a lot of discretion in the hands of the Commission.

To address this problem, the proposed law does two things. First, it encourages the creation of a standard-setting industry body, which “shall develop, on a continuing basis, guidance and bright-line rules for the development and design of technology products of large online operators ....”<sup>166</sup> And second, it directs this industry body to “define conduct that does not have the purpose or substantial effect of subverting or impairing user autonomy, decision-making, or choice ... such as ... de minimis user interface changes derived from testing consumer preferences, including different styles, layouts, or text, where such changes are not done with the purpose of obtaining user consent or user data [and] establishing default settings that provide enhanced privacy protection to users or otherwise enhance their autonomy and decision-making ability.”<sup>167</sup> As this language shows, legislative proponents of clamping down on dark patterns are concerned about the line-drawing problem but feel that without industry input the false-positives problem may become intractable. As our data show, that worry is overblown. Dark patterns were developed through A-B testing, and A-B testing can be used to develop relatively clear and predictable rules about what is permissible. As we explain elsewhere, well-designed surveys are reliable measures for measuring consumer preferences too, so differentiating sludges that seek to undermine widespread preferences from nudges that seek to give consumers what they want is pretty straightforward.<sup>168</sup> An industry association will be more prone to capture than the F.T.C. would be, so there is reason to think that the “more likely than not” test we propose here not only offers a clearer rule but also opens up appealing institutional enforcement options. It could be incorporated into any legislation that tries to address dark patterns.

#### E. Persuasion

A final, tricky, challenge for a systematic effort to regulate dark patterns is to confront the issue of how to deal with variants of dark patterns that may be constitutionally protected.

---

<sup>164</sup> See *supra* text accompanying note 7.

<sup>165</sup> Deceptive Experiences to Online Users Reduction Act (DETOUR Act), Senate Bill 1084 §3(a)(1)(A), 116th Congress, introduced April 9, 2019, text available at <https://www.congress.gov/bill/116th-congress/senate-bill/1084/text>.

<sup>166</sup> *Id.* at § 3(c)(3)(A).

<sup>167</sup> *Id.* at § 3(c)(3)(B)(i)-(iii).

<sup>168</sup> See Strahilevitz & Luguri, *supra* note 25.

For most types of dark patterns, this is relatively easy – false and misleading commercial speech is not protected by the First Amendment.<sup>169</sup> Returning to our taxonomy of dark patterns, then, this means that regulating several categories of dark patterns (social proof, sneaking, forced action, and urgency) is constitutionally unproblematic. In our revised taxonomy we have been more careful than the existing literature to indicate that social proof (activity messages and testimonials) and urgency (low stock / high demand / limited time messages) are only dark patterns insofar as the information conveyed is false or misleading. If a consumer is happy with a product and provides a favorable quote about it, it isn't a dark pattern to use that quote in online marketing, absent a showing that it is misleadingly atypical. Similarly, Amazon can indicate that quantities of an item are limited if there really are unusually low quantities available and if the restocking process could take long enough to delay the customer's order. But the First Amendment's tolerance for the imposition of sanctions on commercial speech is premised on the false character of the defendant's representations, such as by making false representations about the source of content on the defendant's web site.<sup>170</sup> This is an important point, one that the existing writing on dark patterns sometimes misses.

Obstruction and interface interference present marginally harder issues. That said, in a leading case relatively blatant examples of these tactics have been deemed deceptive practices in trade. As such, the conduct would not receive First Amendment protection.<sup>171</sup> But strategies like toying with emotion, as well as confirmshaming, may be hard to restrict under current doctrine given firms' speech interests. There is virtually no legal authority addressing the question of whether commercial speech that satisfies the F.T.C.'s test for unfairness, but is neither misleading nor deceptive, is protected by the First Amendment.<sup>172</sup> The appellate cases that have been litigated recently tend to involve truthful but incomplete disclosures that create a misimpression among consumers, and F.T.C. action in those cases has generally been deemed constitutionally permissible.<sup>173</sup>

Nagging presents perhaps the thorniest type of dark pattern from a First Amendment perspective. CNN's web site employs a nagging dark pattern, one that regularly asks users whether they wish to turn on notifications. There is no question that CNN's core business is protected by the First Amendment. Would a regulation that prevented them from asking consumers to turn on notifications more than once a month, or once a year, infringe on the company's rights as an organization? It would seem not, so long as the rule were implemented as a broadly applicable, content-neutral rule. Here a helpful analogy is to the Federal Do Not Call registry, which applies to newspapers and other speech-oriented entities, but which has withstood First Amendment challenges.<sup>174</sup> Limits on requests to reconsider previous choices

---

<sup>169</sup> *Central Hudson Gas & Electric Corp. v. Public Serv. Comm'n of N.Y.*, 447 U.S. 557, 563 (1980).

<sup>170</sup> *Fanning v. F.T.C.*, 821 F.3d 164, 174-75 (1st Cir. 2016).

<sup>171</sup> See, e.g., *F.T.C. v. AMG Capital Mgmt., LLC*, 910 F.3d 417 (9th Cir. 2018) (involving autorenewal, hidden costs, forced continuity, aesthetic manipulation, and preselection).

<sup>172</sup> See Jennifer L. Pomeranz, *Federal Trade Commission's Authority to Regulate Marketing to Children: Deceptive vs. Unfair Rulemaking*, 21 HEALTH MATRIX 521, 550-52 (2011).

<sup>173</sup> See, e.g., *POM Wonderful LLC v. F.T.C.*, 777 F.3d 478 (D.C. Cir. 2015); *Fanning*, 821 F.3d at 164; *ECM BioFilms, Inc. v. F.T.C.*, 831 F.3d 599 (6th Cir. 2017).

<sup>174</sup> See *Mainstream Marketing Servs. Inc. v. F.T.C.*, 358 F.3d 1228, 1236-46 (10th Cir. 2004) (rejecting a First Amendment challenge to the federal do-not-call registry and holding that the registry's limits on

seem likely to survive similar challenges, provided they establish default rules rather than mandatory ones.<sup>175</sup> On the other hand, the do-not-call cases involve communications by firms to individuals with whom they do not have existing relationships. In the case of nagging restrictions, the government would be limiting what firms can say to their customers in an effort to persuade them to waive existing rights, and it could be that this different dynamic alters the legal bottom line.

Given the potential uncertainty over whether nagging and other forms of annoying-but-nondeceptive forms of dark patterns can be punished, the most sensible strategy for people interested in curtailing these dark patterns is to push on the contractual lever. That is, the First Amendment may be implicated by the imposition of sanctions on firms that nag consumers into agreeing to terms and conditions that do not serve their interests. But there is no First Amendment problem whatsoever with a court or legislature deciding that consent secured via those tactics is voidable. At least in the American legal regime, then, while there is a lot to be gained from considering dark patterns as a key conceptual category, there are some benefits to disaggregation and context-sensitivity, at least in terms of thinking about ideal legal responses.

More broadly, the contractual lever may be the most attractive one for reasons that go far beyond First Amendment doctrine. The F.T.C. has brought some important cases, but neither the federal agency nor enforcers of similar state laws can be everywhere. Public enforcement resources are necessarily finite. But consumers, and attorneys willing to represent them in contract disputes, are numerous. The widespread use of dark patterns could open up firms to substantial class action exposure. As a result, for even a few courts to hold that the use of unfair or deceptive dark patterns obviates consumer consent would significantly deter that kind of conduct.

## Conclusion

Computer scientists discovered dark patterns about a decade ago, and there is a sense in which what they have found is the latest manifestation of something very old – sales practices that test the limits of law and ethics. There is a lot to be learned from looking backwards, but the scale of dark patterns, their rapid proliferation, the possibilities of using algorithms to detect them, and the breadth of the different approaches that have already emerged means this is a realm where significant legal creativity is required.

That is not to say that legal scholars concerned about dark patterns and the harms they can impose on consumers are writing on a blank slate. In a series of unheralded F.T.C. deception cases, and in a few unfairness enforcement actions to boot, the regulator best positioned to address dark patterns has successfully shut down some of the most egregious ones. Courts have generally been sympathetic to these efforts, intuiting the dangers posed by these techniques for consumers' autonomy and their pocketbooks. But an observer of the court cases comes away with an impression that the judges in these cases are like the blind men in the parable of the elephant. They do not understand the interconnectedness of the emerging strategies, nor does the nature of judging allow them to make comparisons about the most pressing problems and

---

telemarketing satisfy the *Central Hudson* test); *National Coalition of Prayer, Inc. v. Carter*, 455 F.3d 783, 787-92 (7th Cir. 2006) (rejecting a First Amendment challenge to a similar state law).

<sup>175</sup> That is, if a customer *wants to be* contacted more than the law provides, they would have the right to permit a commercial speaker to do so. This proviso is important to the constitutional analysis, as *Mainstream Marketing* emphasized that the do not call registry merely established a default.



needs. As a result, they have not given serious thought to the hardest problem facing the legal system – how to differentiate tolerable from intolerable dark patterns.

We think of this paper as making three important contributions to a literature that is growing beyond the human-computer interactions field. First and foremost, there is now an academic paper that demonstrates the effectiveness of various dark patterns. That wasn't true yesterday, even if part of our bottom line is an empirical assessment that has been presupposed by some courts and regarded skeptically by others. The apparent proliferation of dark patterns in ecommerce suggests that they were effective in getting consumers to do things they might not otherwise do, and we now have produced rather solid evidence that this is the case. Paradoxically, it appears that relatively subtle dark patterns are most dangerous, because they sway large numbers of consumers without provoking the level of annoyance that will translate into lost goodwill. Obviously there is a lot more experimental work to do, but this is a critical first step. We hope other social scientists follow us into this body of experimental research.

Second, we have shown how the available experimental evidence helpfully points towards a bright line rule that can be employed to address the aforementioned boundary question. We propose a *per se* rule that treats a dark pattern technique or combination of techniques that more than doubles consumer assent as presumptively unlawful. Our “more likely than not” test is not a panacea – establishing the neutral choice architecture that is to be used as a baseline for comparison is no breeze, and legal judgments about what conduct counts as constitutionally protected “persuasion” must still be made. The *per se* rule will be underinclusive, and it will need to be supplemented by a standard. Yet we think the test we have proposed is workable and desirable.

Third, though legal commentators have largely failed to notice, the F.T.C. is beginning to combat dark patterns with some success, at least in court. The courts are not using the terminology of dark patterns, and they have been hamstrung by the absence of data similar to what we report here. But they have established some key and promising benchmarks already, with the prospect of more good work to come. Developing a systemic understanding of the scope of the problem, the magnitude of the manipulation that is occurring, and the legal landmarks that constrain what the government can do will only aid that new and encouraging effort.

The problem we identify here is both an old problem and a new one. Companies have long manipulated consumers through vivid images, clever turns of phrase, attractive spokesmodels, or pleasant odors and color schemes in stores. This behavior should worry us a little, but not enough to justify aggressive legal responses. Regulating this conduct is expensive, and the techniques are limited in their effectiveness, especially when consumers have the opportunity to learn from previous mistakes.

The online environment is different. It's perhaps only a difference of degree, but the degrees are very large. Through A-B testing, firms now have opportunities to refine and perfect dark patterns that their *Mad Men*-era counterparts could have never imagined. By running tens of thousands of consumers through interfaces that were identical in every respect but one, firms can determine exactly which interface, which text, which juxtapositions, and which graphics maximize revenues. What was once an art is now a science. As a result, consumers' ability to

defend themselves has degraded. The trend toward personalization could make it even easier to weaponize dark patterns against consumers.<sup>176</sup>

Today the law faces a new technology that presents challenges and opportunities. An analogous dynamic has developed recently with partisan gerrymandering and cell tower geolocation. Partisan gerrymandering has been around for a long time, but computing advances in the last several years have made the state-of-the-art techniques precise at a level entirely without precedent, permitting parties to create much greater partisan advantages than they used to be able to. Once the computers became powerful enough, scholars argued that new legal regimes were warranted.<sup>177</sup> But a bitterly divided Supreme Court ultimately disagreed, at least where the federal Constitution is concerned.<sup>178</sup> A similar challenge arose with geolocation, albeit with different results. It had long been settled that police officers could physically tail suspects without a warrant, but when doing just that became trivially expensive, because cell tower records revealed nearly every person's historic whereabouts, scholars said that legal innovation was necessary.<sup>179</sup> And this time the Supreme Court majority agreed with the scholars.<sup>180</sup>

The technology of dark patterns has taken a quantum leap forward, rendering cheap and effective corporate tactics that used to be costly and clunky. So we are making a similar kind of argument to those who suggested that gerrymandering and geolocation technologies had upset status quo assumptions in fundamental ways. Manipulation in the marketplace is a longstanding problem, but recent events have made the problem much worse, and the data presented here gives the strongest hint yet of how large the mismatch is between what consumers want and what they are supposedly consenting to. Dark patterns are a problem that is only going to get worse, because consumers do not have the tools to solve the problem for themselves. Judges, legislators, and regulators now have the data they need to decide whether and how to help.

---

<sup>176</sup> STRAHILEVITZ ET AL., *supra* note 13, at 34-36.

<sup>177</sup> The most prominent such argument is Nicholas O. Stephanopoulos & Eric M. McGhee, *Partisan Gerrymandering and the Efficiency Gap*, 82 U. CHI. L. REV. 831, 868-76, 899-900 (2015).

<sup>178</sup> *Rucho v. Common Cause*, 139 S. Ct. 2484 (2019).

<sup>179</sup> See, e.g., Matthew Tokson, *Knowledge and the Fourth Amendment*, 111 NW. U. L. REV. 139 (2016); Brief of Amici Curiae Empirical Fourth Amendment Scholars in Support of Petitioner, *Carpenter v United States*, No 16-402 (US filed Aug 14, 2017) (available on Westlaw at 2017 WL 3530963) (Strahilevitz signed and was a primary author of that brief).

<sup>180</sup> *Carpenter v. United States*, 138 S Ct 2206 (2018).

Appendix A

Condition	Total	Q1	Q2	Q3	Q4	Q5	Q6	Q7
Description		Accept/Options	Other options	Info 1	Info 2	Info 3	Trick	Reason
<b>Control group</b>	73	73	n/a	n/a	n/a	n/a	n/a	n/a
<b>Mild</b>	155	117	35	n/a	n/a	n/a	n/a	3
<b>Aggressive</b>	217	141	22	11	10	8	24	1

Appendix B

Condition	Overall (% accept)	Low Stakes	High Stakes
<b>Control group</b>	11.3%	9.3%	13.3%
<b>Mild</b>	25.8%	26.8%	24.9%
<b>Aggressive</b>	41.9%	40.9%	42.8%

Appendix C

(When controlling for gender, age, income, race, and education). Numbers represented the beta (standard error).

Trait	Overall	Control group	Mild	Aggressive
<b>Extraversion</b>	.10 (.04) p=.02	.20 (.09) p=.04	.13 (.08) p=.09	.03 (.07) p=.71
<b>Agreeableness</b>	-.09 (.05) p=.08	-.10 (.11) p=.36	-.08 (.08) p=.37	.01 (.09) p=.93
<b>Conscientiousness</b>	-.16 (.05) p=.001	-.29 (.10) p=.006	-.24 (.08) p=.004	-.01 (.08) p=.91
<b>Neuroticism</b>	-.07 (.04) p=.08	-.14 (.09) p=.11	-.04 (.07) p=.54	-.03 (.07) p=.63
<b>Openness</b>	-.05 (.05) p=.29	-.12 (.11) p=.27	-.06 (.08) p=.51	.05 (.08) p=.57