Dark Patterns
Past, Present, and Future

THE EVOLUTION OF TRICKY USER INTERFACES

ARVIND NARAYANAN, ARUNESH MATHUR, MARSHINI CHETTY, AND MIHIR KSHIRSAGAR

Dark patterns are user interfaces that benefit an online service by leading users into making decisions they might not otherwise make. Some dark patterns deceive users while others covertly manipulate or coerce them into choices that are not in their best interests. A few egregious examples have led to public backlash recently: TurboTax hid its U.S. government-mandated free tax-file program for low-income users on its website to get them to use its paid program;[9] Facebook asked users to enter phone numbers for two-factor authentication but then used those numbers to serve targeted ads;[31] Match.com knowingly let scammers generate fake messages of interest in its online dating app to get users to sign up for its paid service.[13] Many dark patterns have been adopted on a large scale across the web. Figure 1 shows a deceptive countdown timer dark pattern on JustFab. The advertised offer remains valid even after the timer expires. This pattern is a common tactic—a recent study found such deceptive countdown timers on 140 shopping websites.[20]

The research community has taken note. Recent efforts

1

FIGURE 1: **A DECEPTIVE COUNTDOWN TIMER ON JUSTFAB**

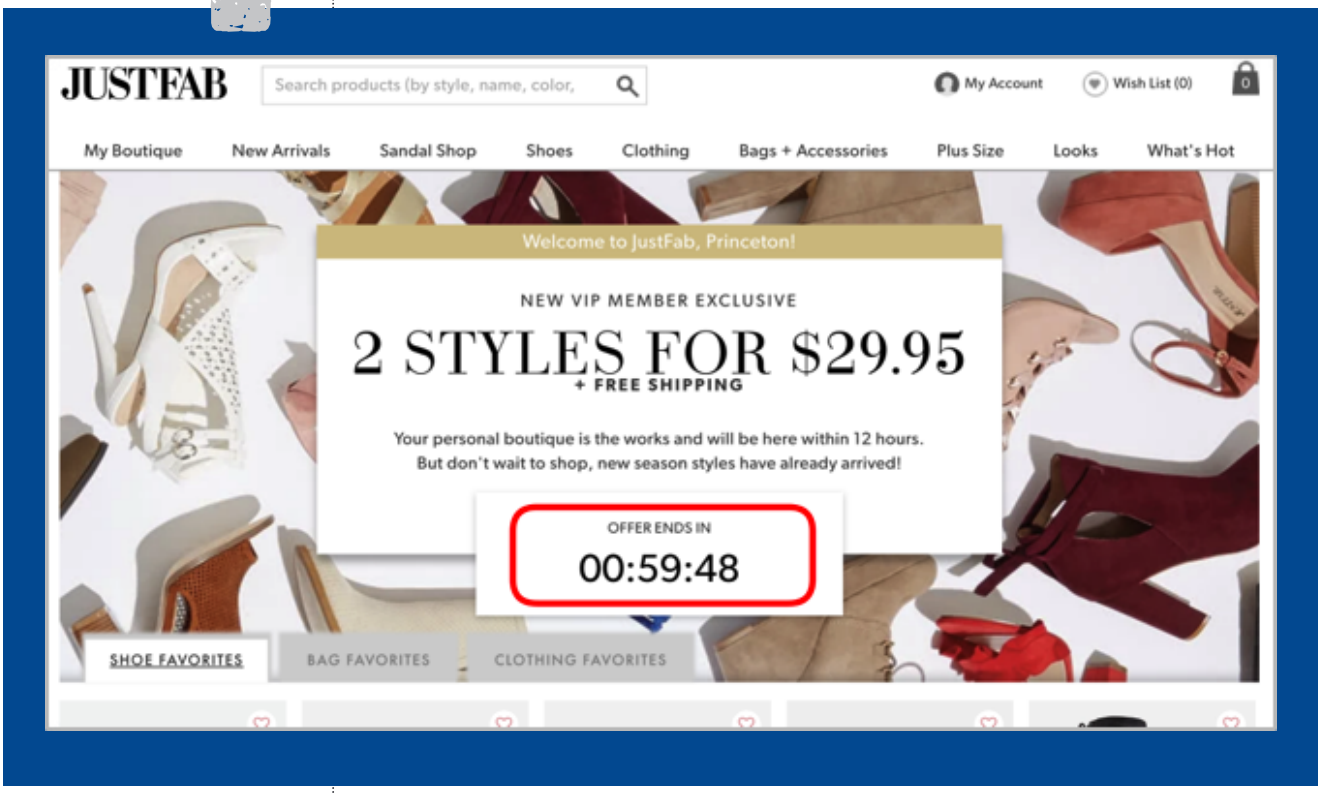

have catalogued dozens of problematic patterns such as nagging the user, obstructing the flow of a task, and setting privacy-intrusive defaults,[1,18] building on an early effort by Harry Brignull (darkpatterns.org). Researchers have also explained how dark patterns operate by exploiting cognitive biases,[4,20,33] uncovered dark patterns on more than 1,200 shopping websites,[20] shown that more than 95 percent of the popular Android apps contain dark patterns,[8] and provided preliminary evidence that dark patterns are indeed effective at manipulating user behavior.[19,30]

Although they have recently burst into mainstream awareness, dark patterns are the result of three decades-long trends: one from the world of retail (deceptive

practices), one from research and public policy (nudging), and the third from the design community (growth hacking).

Figure 2 illustrates how dark patterns stand at the confluence of these three trends. Understanding these trends—and how they have collided into each other—is essential to help us appreciate what is actually new about dark patterns, demystifies their surprising effectiveness, and shows us why it will be hard to combat them. We end this article with recommendations for ethically minded designers.

DECEPTION AND MANIPULATION IN RETAIL

The retail industry has a long history of deceptive and manipulative practices that range on a spectrum from normalized to unlawful (figure 3). Some of these techniques, such as psychological pricing (i.e., making the price slightly less than a round number), have become

2

FIGURE 2: **THE ORIGINS OF DARK PATTERNS**

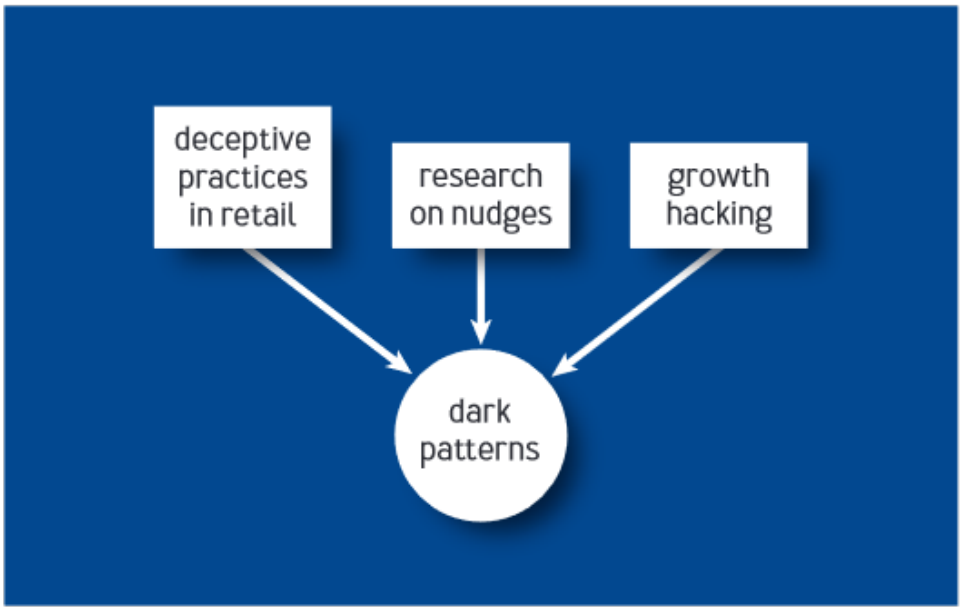

FIGURE 3: **EXAMPLES OF DECEPTIVE AND MANIPULATIVE RETAIL PRACTICES**

(a) psychological pricing

Source: https://www.crazyspeedtech.com/5-major-stages-psychological-pricing/

(b) false advertisement of store closure

Source: https://www.dealnews.com/features/What-Happens-When-a-Store-Closes/2203265.html

(c) bait-and-switch car ad

Source: https://www.ftc.gov/enforcement/cases-proceedings/1223269/ganley-ford-west-inc-matter

normalized. This is perfectly legal, and consumers have begrudgingly accepted it. Nonetheless, it remains effective: consumers underestimate prices when relying on memory if psychological pricing is employed.[3]

More problematic are practices such as false claims of store closings, which are unlawful but rarely the target of enforcement actions. At the other extreme are bait-and-switch car ads such as the one by a Ford dealership in Cleveland that was the target of an FTC action.[14]

THE ORIGINS OF NUDGING

In the 1970s, the heuristics and biases literature in behavioral economics sought to understand irrational decisions and behaviors—for example, people who decide to drive because they perceive air travel as dangerous, even though driving is, in fact, orders of magnitude more dangerous per mile.[29] Researchers uncovered a set of cognitive shortcuts used by people that make these irrational behaviors not just explainable but even predictable.

For example, in one experiment, researchers asked participants to write down an essentially random two-digit number (the last two digits of each participant's social security number), then asked if they would pay that number of dollars for a bottle of wine, and finally asked the participants to state the maximum amount they would pay for the bottle.[2] They found that the willingness to pay varied by roughly threefold based on the arbitrary number. This is the *anchoring effect*: lacking knowledge of the market value of the bottle of wine, participants' estimates become anchored to the arbitrary reference point. This

study makes it easy to see how businesses might be able to nudge customers to pay higher prices by anchoring their expectations to a high number. In general, however, research on psychological biases has not been driven by applications in retail or marketing. That would come later.

NUDGING: THE TURN TO PATERNALISM

The early behavioral research on this topic focused on understanding rather than intervention. Some scholars, such as Cass Sunstein and Richard Thaler, authors of the book *Nudge,*[28] went further to make a policy argument: Governments, employers, and other benevolent institutions should engineer "choice architectures" in a way that uses behavioral science for the benefit of those whom they serve or employ.

A famous example (figure 4) is the striking difference in organ-donation consent rates between countries where people have to explicitly provide consent (red bars) versus those where consent is presumed (orange bars). Because most people tend not to change the default option, the latter leads to significantly higher consent rates.[17]

Today, nudging has been enthusiastically adopted by not only governments and employers, but also businesses in the way they communicate with their customers. The towel reuse message you may have seen in hotel rooms ("75 percent of guests in this hotel usually use their towels more than once") is effective because it employs descriptive social norms as a prescriptive rule to get people to change their behavior.[16]

With the benefit of hindsight, neither the proponents nor the critics of nudging anticipated how readily and

4

FIGURE 4: **ORGAN-DONATION CONSENT RATES BY COUNTRIES**

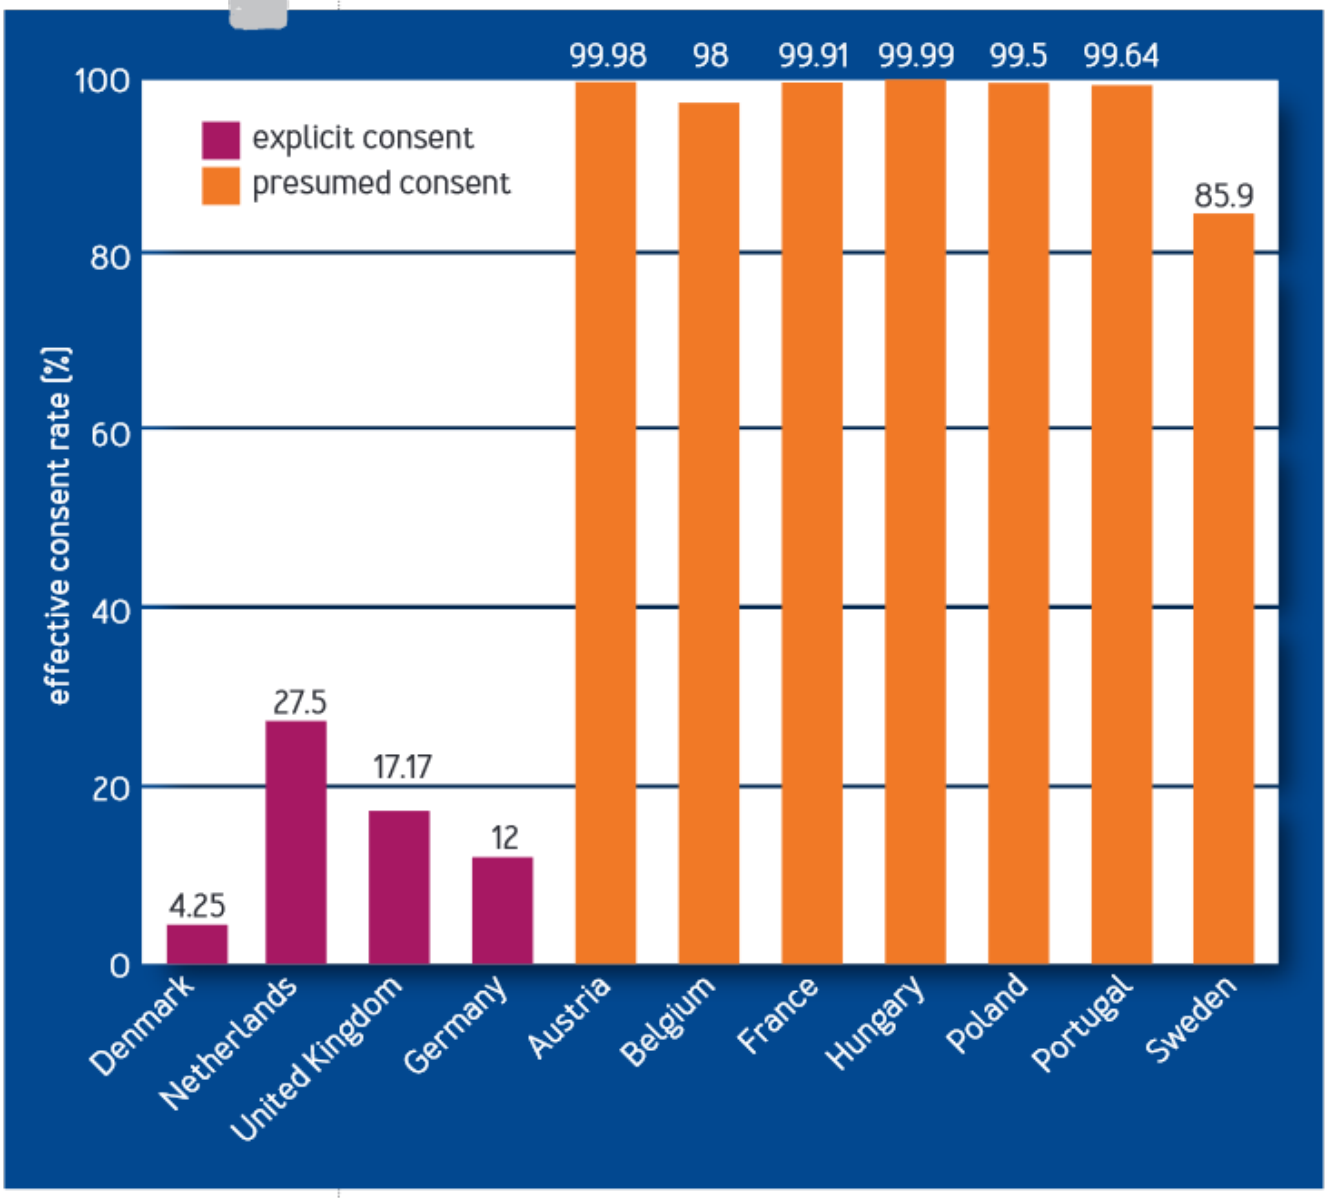

vigorously businesses would adopt these techniques in adversarial rather than paternalistic ways. In *Nudge* Sunstein and Thaler briefly address the question of how to tell if a nudge is ethical, but the discussion is perfunctory. The authors seem genuinely surprised by recent developments and have distanced themselves

from dark patterns, which they label "sludges."[27]

GROWTH HACKING

The third trend—and the one that most directly evolved into dark patterns—is growth hacking. The best-known and arguably the earliest growth hack was implemented by Hotmail. When it launched in 1996, the founders first considered traditional marketing methods such as billboard advertising. Instead, they hit upon a viral marketing strategy: The service automatically added the signature, "Get your free email with Hotmail," to every outgoing email, essentially getting users to advertise on its behalf, resulting in viral growth.[21]

Successes like these led to the emergence of growth hacking as a distinct community. Growth hackers are trained in design, programming, and marketing and use these skills to drive product adoption.

Growth hacking is not inherently deceptive or manipulative but often is in practice. For example, in two-sided markets such as vacation rentals, upstarts inevitably face a chicken-and-egg problem: no travelers without hosts and no hosts without travelers. So it became a common practice to "seed" such services with listings that were either fake or scraped from a competitor.[22,23]

Unsurprisingly, growth hacking has sometimes led to legal trouble. A hugely popular growth hack involved obtaining access to users' contact books—often using deception—and then spamming those contacts with invitations to try a service. The invitations might themselves be deceptive by appearing to originate from the user, when in fact users were unaware of the emails

being sent. Linkedin settled a class action for exactly this practice, which it used from 2011 to 2014.[25]

FROM GROWTH HACKING TO DARK PATTERNS

But why growth rather than revenue or some other goal? It is a reflection of Silicon Valley's growth-first mantra in which revenue-generating activities are put aside until after-market dominance has been achieved. Of course, eventually every service runs into limits on growth, because of either saturation or competition, so growth hackers began to adapt their often-manipulative techniques to extracting and maximizing revenue from existing users.

In developing their battery of psychological tricks, growth hackers had two weapons that were not traditionally available in offline retail. The first was that the nudge movement had helped uncover the principles of behavior change. In contrast, the marketing literature that directly studied the impact of psychological tricks on sales was relatively limited because it didn't get at the foundational principles and was limited to the domain of retail.

The second weapon was A/B testing (figure 5). By serving variants of web pages to two or more randomly

FIGURE 5: **HYPOTHETICAL ILLUSTRATION OF A/B TESTING ON A WEBSITE**

Source: https://en.wikipedia.org/wiki/A/B_testing

selected subsets of users, designers began to discover that even seemingly trivial changes to design elements can result in substantial differences in behavior. The idea of data-driven optimization of user interfaces has become deeply ingrained in the design process of many companies. For large online services with millions of users, it is typical to have dozens of A/B tests running in parallel, as noted in 2009 by Douglas Bowman, once a top visual designer at Google:

> *Yes, it's true that a team at Google couldn't decide between two blues, so they're testing 41 shades between each blue to see which one performs better. I had a recent debate over whether a border should be 3, 4, or 5 pixels wide, and was asked to prove my case. I can't operate in an environment like that. I've grown tired of debating such minuscule design decisions. There are more exciting design problems in this world to tackle.*
>
> —Douglas Bowman

A/B testing proved key to the development of dark patterns because it is far from obvious how to translate an abstract principle like social proof into a concrete nudge ("7 people are looking at this hotel right now!"). Another example: For how long should a fake countdown timer be set ("This deal expires in 15 minutes!" ... "14:59" ... "14:58" ...), so that the user acts with urgency but not panic? Online experiments allow designers to find the answers with just a few lines of code.

MONEY, DATA, ATTENTION

Let's recap. As the online economy matured, services turned their attention from growth to revenue. They used the principles of behavioral influence but subverted the intent of the researchers who discovered those principles by using them in ways that undermined consumers' autonomy and informed choice. They used A/B testing to turn behavioral insights into strikingly effective user interfaces. In some cases these were optimized versions of tricks that have long been used in retail, but in other cases they were entirely new.

How, exactly, do dark patterns help maximize a company's ability to extract revenue from its users? The most obvious way is simply to nudge (or trick) consumers into spending more than they otherwise would.

A less obvious, yet equally pervasive, goal of dark patterns is to invade privacy. For example, cookie consent dialogs almost universally employ manipulative design to increase the likelihood of users consenting to tracking. In fact, a recent paper shows that when asked to opt in, well under 1 percent of users would provide informed consent.[30] Regulations such as the GDPR (General Data Protection Regulation) require companies to get explicit consent for tracking, which poses an existential threat to many companies in the online tracking and advertising industry. In response, they appear to be turning to the wholesale use of dark patterns.[30]

A third goal of dark patterns is to make services addictive. This goal supports the other two, as users who stay on an app longer will buy more, yield more personal information, and see more ads. Apps like Uber use gamified

FIGURE 6: **ONE OF UBER'S GAMIFIED NUDGES TO KEEP DRIVERS ON THE ROAD**

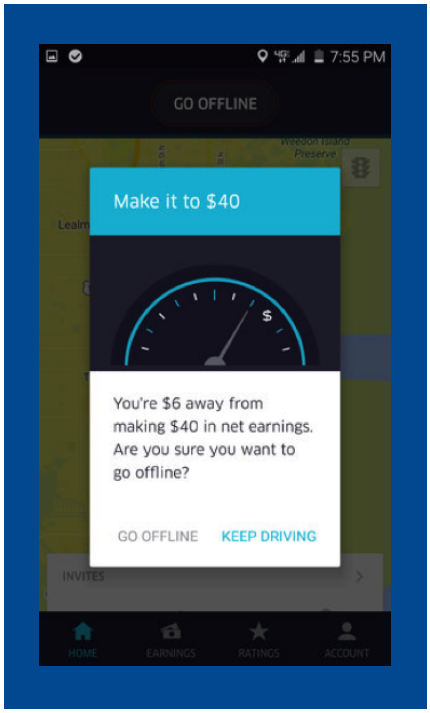

Source: https://www.nytimes.com/interactive/2017/04/02/technology/uber-drivers-psychological-tricks.html

nudges to keep drivers on the road longer (figure 6). The needle suggests that the driver is extremely close to the goal, but it is an arbitrary goal set by Uber when a driver wants to go offline.[24] To summarize, dark patterns enable designers to extract three main resources from users: money, data, and attention.

DARK PATTERNS ARE HERE TO STAY

Two years ago, few people had heard the term *dark patterns*. Now it's everywhere. Does this mean dark patterns are a flash in the pan? Perhaps, as users figure out what's going on, companies will realize that dark patterns are counterproductive and stop using them. The market could correct itself.

The history sketched here suggests that this optimistic

view is unlikely. The antecedents of dark patterns are decades old. While public awareness of dark patterns is relatively new, the phenomenon itself has developed gradually. In fact, the darkpatterns.org website was established in 2010.

The history also helps explain what is new about dark patterns. It isn't just tricky design or deceptive retail practices online. Rather, design has been weaponized using behavioral research to serve the aims of the surveillance economy. This broader context is important. It helps explain why the situation is as bad as it is and suggests that things will get worse before they can get better.

One worrying trend is the emergence of companies that offer dark patterns as a service, enabling websites to adopt them with a few lines of JavaScript.[20] Another possible turn for the worse is personalized dark patterns that push each user's specific buttons.[26] This has long been predicted[5] but remains rare today (manipulative targeted advertising can arguably be viewed as a dark pattern, but ads are not user interfaces). The absence of personalized UI is presumably because companies are busy picking lower-hanging fruit, but this can change any time.

RECOMMENDATIONS FOR DESIGNERS

Designers should be concerned about the proliferation of dark patterns. They are unethical and reflect badly on the profession. But this article is not a doom-and-gloom story. There are steps you can take, both to hold yourself and your organization to a higher standard, and to push back against the pressure to deploy dark patterns in the industry.

Go beyond superficial A/B testing metrics

Earlier this article discussed how designers use A/B tests to optimize dark patterns. But there's a twist: a design process hyperfocused on A/B testing can result in dark patterns even if that's not the intent. That's because most A/B tests are based on metrics that are relevant to the company's bottom line, even if they result in harm to users. As a trivial example, an A/B test might reveal that reducing the size of a "Sponsored" label that identifies a search result as an advertisement causes an increase in the CTR (click-through rate). While a metric such as CTR can be measured instantaneously, it reveals nothing about the long-term effects of the design change. It is possible that users lose trust in the system over time when they realize they are being manipulated into clicking on ads.

In fact, Google's recent change to its ad labels made it hard for users to distinguish ads from organic search results, and presumably increased CTR for ads (figure 7). A backlash ensued, however, and Google rolled back this interface.[32]

To avoid falling into this trap, evaluate A/B tests on at least one metric that measures long-term impacts. In addition to measuring the CTR, you could also measure user retention. That will tell you if a different-sized label results in more users abandoning the website.

Still, many attributes that matter in the long term, such as trust, are not straightforward to observe and measure, especially in the online context. Think critically about the designs you choose to test, and when you find that a certain design performs better, try to understand why.

While the overreliance on A/B testing is a critical issue

FIGURE 7: **GOOGLE'S RECENT CHANGE TO ITS AD LABELS**

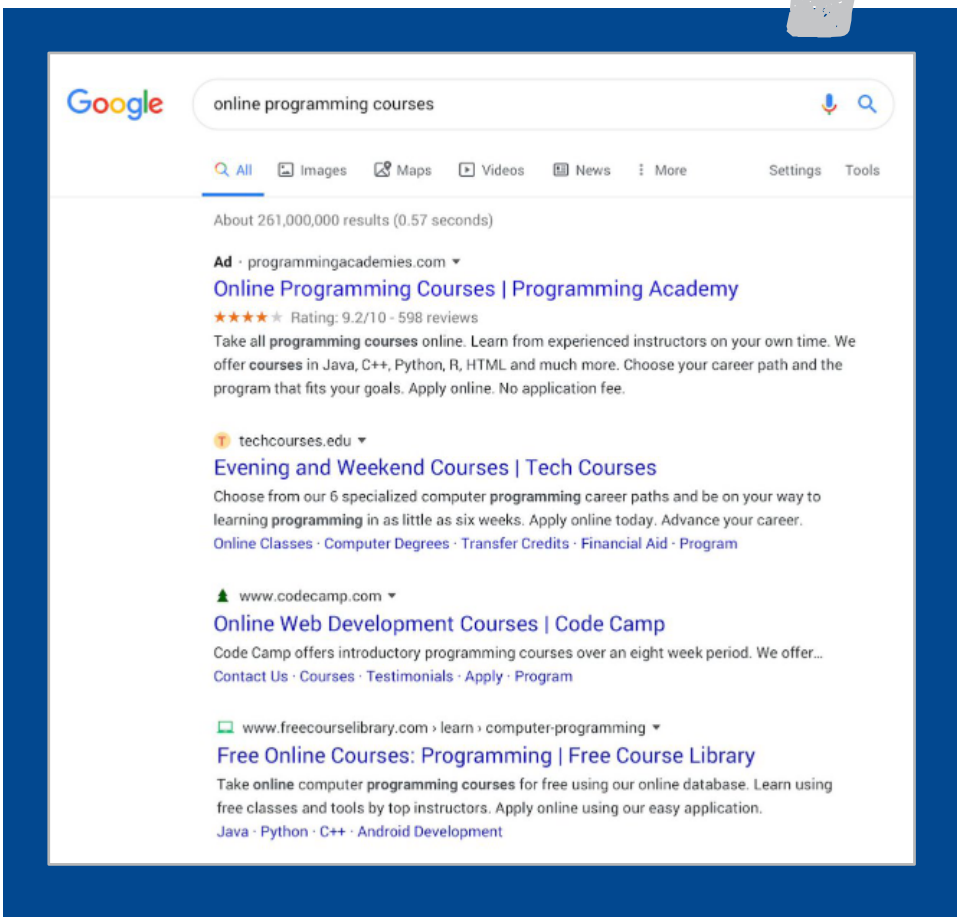

to be addressed, let's next turn to a much broader and longer-term concern.

Incorporate ethics into the design process

While dark patterns are a highly visible consequence of the ethical crisis in design, resolving the crisis entails far more than avoiding a simple list of patterns. It requires structural changes to the design process.

Start by articulating the values that matter to you and that will guide your design.[15] Not every organization will have an identical set of values, but these values must be

broadly aligned with what society considers important.

In fact, much of the present crisis can be traced to a misalignment of values between society and companies. Autonomy and privacy are two values where this is particularly stark. Consider frictionless design, a bedrock value in the tech industry. Unfortunately, it robs users of precisely those moments that may give them opportunities for reflection and enable them to reject their baser impulses. Frictionlessness is antithetical to autonomy. Similarly, designing for pleasure and fun is a common design value, but when does fun cross the line into addiction?

Once you've articulated your values, continue to debate them internally. Publicize them externally, seek input from users, and, most importantly, hold yourself accountable to them. Effective accountability is challenging, however. For example, advisory boards established by technology companies have been criticized for not being sufficiently independent.

Everyday design decisions should be guided by referring to established values. In many cases it is intuitively obvious whether a design choice does or does not conform to a design value, but this is not always so. Fortunately, research has revealed a lot about the factors that make a design pattern dark, such as exploiting known cognitive biases and withholding crucial information.[4,20] Stay abreast of this research, evaluate the impact of design on your users, and engage in critical debate about where to draw the line based on the company's values and your own sense of ethics. Rolling back a change should always be an option if it turns out that it didn't live up to your values.

As you gain experience making these decisions in a particular context, higher-level principles can be codified into design guidelines. There is a long tradition of usability guidelines in the design community. There are also privacy-by-design guidelines, but they are not yet widely adopted.[10] There is relatively little in the way of guidelines for respecting user autonomy.

All of this is beyond the scope of what individual designers can usually accomplish; the responsibility for incorporating ethics into the design process rests with organizations. As an individual, you can start by raising awareness within your organization.

Self regulate or get regulated

Dark patterns are an abuse of the tremendous power that designers hold in their hands. As public awareness of dark patterns grows, so does the potential fallout. Journalists and academics have been scrutinizing dark patterns, and the backlash from these exposés can destroy brand reputations and bring companies under the lenses of regulators.

Many dark patterns are already unlawful. In the United States, the Federal Trade Commission (FTC) Act prohibits "unfair or deceptive" commercial practices.[11] In a recent example, the FTC reached a settlement with Unroll.Me—a service that unsubscribed users' email addresses from newsletters and subscriptions—because it was in fact selling information it read from their inboxes to third parties.[12] European Union authorities have tended to be stricter: French regulator CNIL (Commission Nationale de l'Informatique et des Libertés) fined Google 50 million

Related articles


User Interface Designers, Slaves of Fashion
The status quo prevails in interface design, and the flawed concept of cut-and-paste is a perfect example.
Jef Raskin
https://queue.acm.org/detail.cfm?id=945161


The Case Against Data Lock-in
Want to keep your users? Just make it easy for them to leave.
Brian W. Fitzpatrick and J.J. Lueck
https://queue.acm.org/detail.cfm?id=1868432


Bitcoin's Academic Pedigree
The concept of cryptocurrencies is built from forgotten ideas in research literature.
Arvind Narayanan and Jeremy Clark
https://queue.acm.org/detail.cfm?id=3136559

euros for hiding important information about privacy and ad personalization behind five to six screens.[6]

There is also a growing sense that existing regulation isn't enough, and new legislative proposals aim to curb dark patterns.[7] While policymakers should act—whether by introducing new laws or by broadening and strengthening the enforcement of existing ones—relying on regulation isn't sufficient and comes with compliance burdens.

Let's urge the design community to set standards for itself, both to avoid onerous regulation and because it's the right thing to do. A first step would be to rectify the misalignment of values between the industry and society, and develop guidelines for ethical design. It may also be valuable to partner with neutral third-party consumer advocacy agencies to develop processes to certify apps that are free of known dark patterns. Self-regulation also requires cultural change. When hiring designers, ask about the ethics of their past work. Similarly, when deciding between jobs, use design ethics as one criterion for evaluating a company and the quality of its work environment.

Design is power. In the past decade, software engineers have had to confront the fact that the power they hold comes with responsibilities to users and to society. In this decade, it is time for designers to learn this lesson as well.

References

1. Acquisti, A., Adjerid, I., Balebako, R., Brandimarte, L., Cranor, L. F., Komanduri, S., Leon, P. G., Sadeh, N., Schaub, F., Sleeper, M., Wang, Y., Wilson, S. 2017. Nudges for privacy and security: understanding and assisting users' choices online. *ACM Computing Surveys* 50(3), 1-41; https://dl.acm.org/doi/10.1145/3054926.
2. Ariely, D. 2008. *Predictably Irrational*. New York, NY: Harper Audio.
3. Bizer, G. Y., Schindler, R. M. 2005. Direct evidence of ending-digit drop-off in price information processing. *Psychology & Marketing* 22(10), 771-783.
4. Bösch, C., Erb, B., Kargl, F., Kopp, H., Pfattheicher, S. 2016. Tales from the dark side: privacy dark strategies and privacy dark patterns. *Proceedings on Privacy Enhancing Technologies* 2016(4), 237-254.
5. Calo, R. 2014. Digital market manipulation. *George Washington Law Review* 82(4), 995-1051; http://www.gwlr.org/wp-content/uploads/2014/10/Calo_82_41.pdf.
6. Commission Nationale de l'Informatique et des Libertés. 2019. The CNIL's restricted committee imposes a financial penalty of 50 million euros against Google LLC; https://www.cnil.fr/en/cnils-restricted-committee-imposes-financial-penalty-50-million-euros-against-google-llc.
7. Deb Fischer, United States Senator for Nebraska.

2019. Senators introduce bipartisan legislation to ban manipulative dark patterns; https://www.fischer.senate.gov/public/index.cfm/2019/4/senators-introduce-bipartisan-legislation-to-ban-manipulative-dark-patterns.

8. Di Geronimo, L., Braz, L., Fregnan, E., Palomba F., Bachelli, A. 2020. UI dark patterns and where to find them: a study on mobile applications and user perception. In *Proceedings of the 2020 ACM Conference on Human Factors in Computing Systems*.
9. Elliott, J., Waldron, L. 2019. Here's how TurboTax just tricked you into paying to file your taxes. *ProPublica* (April 22); https://www.propublica.org/article/turbotax-just-tricked-you-into-paying-to-file-your-taxes.
10. European Data Protection Board. 2019. Guidelines 4/2019 on Article 25, Data Protection by Design and by Default; https://edpb.europa.eu/sites/edpb/files/consultation/edpb_guidelines_201904_dataprotection_by_design_and_by_default.pdf.
11. Federal Trade Commission. 2019. A brief overview of the Federal Trade Commission's investigative, law enforcement, and rulemaking authority; https://www.ftc.gov/about-ftc/what-we-do/enforcement-authority.
12. Federal Trade Commission. 2019. FTC finalizes settlement with company that misled consumers about how it accesses and uses their email; https://www.ftc.gov/news-events/press-releases/2019/12/ftc-finalizes-settlement-company-misled-consumers-about-how-it.
13. Federal Trade Commission. 2019. FTC sues owner of online dating service Match.com for using fake love interest ads to trick consumers into paying for a Match.

com subscription. Sept. 25; https://www.ftc.gov/news-events/press-releases/2019/09/ftc-sues-owner-online-dating-service-matchcom-using-fake-love.

14. Federal Trade Commission. Ganley Ford; https://www.ftc.gov/sites/default/files/documents/cases/2013/09/130903ganleyfordexhibita.pdf.
15. Friedman, B., Kahn, P. H., Borning, A., Huldtgren, A. 2013. Value-sensitive design and information systems. In *Early Engagement and New Technologies: Opening up the Laboratory*, ed. N. Doorn, D. Schuurbiers, I. van de Poel, M. E. Gorman, 55-95. Dordrecht: Springer; https://link.springer.com/book/10.1007/978-94-007-7844-3.
16. Goldstein, N. J., Cialdini, R. B., Griskevicius, V. 2008. A room with a viewpoint: using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research* 35(3), 472-482.
17. Goldstein, D., Johnson, E. J. 2003. Do defaults save lives? *Science* 302(5649), 1338-1339; https://science.sciencemag.org/content/302/5649/1338.
18. Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., Toombs, A. L. 2018. The dark (patterns) side of UX design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (April), 1-14; https://dl.acm.org/doi/10.1145/3173574.3174108.
19. Luguri, J., Strahilevitz, L. 2019. Shining a light on dark patterns. University of Chicago, Public Law Working Paper. No. 719.
20. Mathur, A., Acar, G., Friedman, M. J., Lucherini, E., Mayer, J., Chetty, M., Narayanan, A. 2019. Dark patterns at scale: findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer*

Interaction 3(CSCW), 1-32; https://dl.acm.org/doi/10.1145/3359183.

21. McLaughlin, J. 2014. 9 iconic growth hacks tech companies used to boost their user bases. *The Next Web*; https://thenextweb.com/entrepreneur/2014/05/28/9-iconic-growth-hacks-tech-companies-used-pump-user-base.
22. Mead, D. 2012. How Reddit got huge: tons of fake accounts. *Vice*; https://www.vice.com/en_us/article/z4444w/how-reddit-got-huge-tons-of-fake-accounts--2.
23. Rosoff, M. 2011. Airbnb farmed Craigslist to grow its listings, says competitor. *Business Insider*; https://www.businessinsider.com/airbnb-harvested-craigslist-to-grow-its-listings-says-competitor-2011-5.
24. Scheiber, N. 2017. How Uber uses psychological tricks to push its drivers' buttons. *New York Times* (April 22); https://www.nytimes.com/interactive/2017/04/02/technology/uber-drivers-psychological-tricks.html.
25. Strange, A. 2015. LinkedIn pays big after class action lawsuit over user emails. *Mashable*; https://mashable.com/2015/10/03/linkedin-class-action.
26. Susser, D., Roessler, B., Nissenbaum, H. 2019. Online manipulation: hidden influences in a digital world. *Georgetown Law Technology Review* 4.1, 1-45; https://philarchive.org/archive/SUSOMHv1.
27. Thaler, R. H. 2018. Nudge, not sludge. *Science* 361, 431-431.
28. Thaler, R. H., Sunstein, C. R. 2009. Nudge: Improving Decisions About Health, Wealth, and Happiness. Penguin Books.

29. Tversky, A., Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.
30. Utz, C., Degeling, M., Fahl, S., Schaub, F., Holz, T. 2019. (Un)informed consent: studying GDPR consent notices in the field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 973-990; https://dl.acm.org/doi/10.1145/3319535.3354212.
31. Venkatadri, G., Lucherini, E., Sapiezynski, P., Mislove, A. 2019. Investigating sources of PII used in Facebook's targeted advertising. *Proceedings on Privacy Enhancing Technologies* 2019(1), 227-244; https://content.sciendo.com/view/journals/popets/2019/1/article-p227.xml?lang=en.
32. Wakabayashi, D., Hsu, T. 2020. Why Google backtracked on its new search results look. *New York Times* (January 31); https://www.nytimes.com/2020/01/31/technology/google-search-results.html.
33. Waldman, A. E. 2019. Cognitive biases, dark patterns, and the 'privacy paradox.' SSRN.

Arvind Narayanan *is an associate professor of computer science at Princeton University. He leads the Princeton Web Transparency and Accountability Project to uncover how companies collect and use our personal information. Narayanan co-created a Massive Open Online Course and textbook on bitcoin and cryptocurrency technologies, which has been used in more than 150 courses worldwide. His recent work has shown how machine learning reflects cultural stereotypes, and his doctoral research showed*

the fundamental limits of de-identification. Narayanan is a recipient of the Presidential Early Career Award for Scientists and Engineers, twice recipient of the Privacy Enhancing Technologies Award, and thrice recipient of the Privacy Papers for Policy Makers Award.

Arunesh Mathur *is a graduate student in the department of computer science at Princeton University. Mathur studies the societal impacts of technical systems through an empirical lens. His recent research examines how commercial, political, and other powerful actors employ dark patterns to exploit individuals and society. His research has received several awards including the Best Paper Award at ACM CSCW (Computer-supported Cooperative Work) and Usenix SOUPS (Symposium on Usable Privacy and Security), and the Privacy Papers for Policy Makers Award.*

Marshini Chetty *is an assistant professor in the department of computer science at the University of Chicago. She specializes in human-computer interaction, usable privacy and security, ubiquitous computing, and inclusive technology. Marshini has a Ph.D. in human-centered computing from the Georgia Institute of Technology and master's and bachelor's degrees in computer science from the University of Cape Town, South Africa. Her work has won best paper awards at SOUPS (Symposium on Usable Privacy and Security), CHI (Conference on Human Factors in Computing Systems), and CSCW (Computer-supported Cooperative Work) and has been funded by the National Science Foundation, the National*

Security Agency, Intel, Microsoft, Facebook, and multiple Google Faculty Research Awards.

Mihir Kshirsagar *leads the Tech Policy Clinic at Princeton CITP (Center for Information Technology Policy). He previously worked for the New York attorney general and was in private practice before that. He has a law degree from the University of Pennsylvania and an undergraduate degree from Harvard.*

CONTENTS