May 1, 2022

Re: EDPB Guidelines 3/2022 on Dark patterns in social media platform interfaces: How to recognise and avoid them

Dear Members of the European Data Protection Board,

We are a group of University of Chicago faculty and students at the University of Chicago Law School, Harris School of Public Policy, and Computer Science Department with areas of expertise in American law on dark patterns, the public policy of coercive digital interfaces, and user-focused, privacy-protective website design. We write to comment on the 3/2022 Dark patterns guidelines adopted on March 14. Given the extraterritorial impact EDPB guidelines often have and the influence of European law on American law and design best practices, these guidelines will be central to our professional and academic lives. As such, we sincerely appreciate the EDPB's commitment to public engagement and the opportunity to share our collective expertise and suggestions. It is our hope that the 3/2022 guidelines grow into a coherent set of best practices applicable in Europe but also to the professional and academic contexts outside of Europe with which we interact on a day-to-day basis. To that end, our comments typically pertain to making the guidelines more accessible to a variety of audiences.

The following text is divided into two sections. First, we present general comments on the accessibility of the guidelines. While it is our hope that the general comments are taken into consideration, in an effort to offer implementable, concrete suggestions, we present some specific feedback on individual sections, paragraphs, and wording in the second section. Though we understand that the EDPB's commitment to collaborative governance over enforcement eases concerns about vagueness, we believe that changing certain wordings and in some cases modifying the examples used in certain use cases will improve readability and persuasiveness.

§ 1: General Comments

At a high level, we commend the EDPB on its commitment to addressing dark patterns. The scholarly consensus is that dark patterns deprive users of autonomy in digital space[1] while appropriating economic value for firms in an anticompetitive manner.[2] They also evince a fundamental lack of respect

[1] *See* Daniel Susser, Beate Roessler & Helen Nissenbaum, *Online Manipulation: Hidden Influences in a Digital World*, 4 GEO. L. TECH. REV. 1, 34-44 (2019). *See also* Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 955 (2014).
[2] Gregory Day & Abbey Stemler, *Are Dark Patterns Anticompetitive?*, 72 ALA. L. REV. 1 (2020) (arguing forcefully that the market power coercive data practices and dark patterns allow firms to accumulate negatively impacts consumer welfare by decreasing market competition).

for individual dignity and privacy that is deeply concerning in an era of "surveillance capitalism."[3] Given this background, we found it somewhat surprising that the guidelines are meant to guide both users and interface designers. While it is certainly true that both groups share a need for information about dark patterns, their often conflicting economic incentives render their information needs of a different kind. Bluntly, designers want clear, easily digestible rules of thumb, not examples of dark patterns. Users, on the other hand, are likely to want a quick set of examples that can help them to identify dark patterns in their everyday lives. In short, designers may want rules and users may want examples. We believe the desire to speak to multiple audiences at once obfuscates the impact of the guidelines, especially since the "best practices" sections are scattered throughout the 64 pages rather than consolidated in a single location. In the future, we would suggest both consolidation of best practices in one location and writing separate guidelines for users and designers.

More broadly, we'd like to point out that not all designers and without question not all users (particularly those outside Europe) will be familiar with definitions of "personal data" ostensibly borrowed from Article 4 of the GDPR and "social media" borrowed from Article 2 of the Digital Services Act. The importance and scope of the guidelines for such individuals thus could be clarified. Outlining–even cursorily–how "personal data" is defined for purposes of the guidelines and what is necessary for a platform or user interface to fall under the heading "social media" would improve the useability of the guidelines outside Europe. The definition of "dark patterns" might also be clarified to improve useability. For example, what counts as a discrete "user experience" is left unspecified and its relationship to "interfaces" is unclear to the non-technical user. Is "user experience" meant to indicate something outside interaction with a website's interface? Are the two categories separate from one another? We believe additional guidance on these points would be very beneficial.

Moreover, a one-size-fits-all definition of "dark patterns" is likely to ignore the empirical fact that not all dark patterns are equally harmful because not all dark patterns are equally effective.[4] Empirically, not all dark patterns harm consumers in the same way or to the same degree and it may be more efficacious for time-pressed designers and users to be encouraged to avoid the most egregious versions of the category.[5] One way to do this would be to consider adding gradations to the present definition of dark patterns based on harm.[6] Separately, a focus on the fact that harm varies would improve the

---

[3] JULIE COHEN, BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM 48-75 (2019). Conversely, for a consequentialist assessment of the argument for reducing manipulation by reducing manipulation online, *see* Cass Sunstein, *Sludge and Ordeals*, 68 DUKE L. J. 1843 (2019).
[4] Jamie Luguri & Lior J. Strahilevitz, *Shining a Light on Dark Patterns*, 13 J. OF LEG. ANAL. 43 (2021).
[5] *Id.*
[6] A four-tiered definition might look something like:

> Category 1 (most harmful): interfaces and user experiences that frequently lead users into making decisions they would not otherwise make and which result in the probable release of sensitive personal data
> Category 2 (harmful): interfaces and user experiences that frequently lead users into making decisions they would not otherwise make and which result in the probable release of non-sensitive personal data
> Category 3 (somewhat harmful): interfaces and user experiences that occasionally but infrequently lead users into making decisions they would not otherwise make and which result in the probable release of sensitive personal data
> Category 4 (least harmful): interfaces and user experiences that occasionally but infrequently lead users into making decisions they would not otherwise make and which result in the probable release of non-sensitive personal data

persuasiveness of the guidelines by countering the tendency among regulated parties to imagine only the least harmful cases in the process of dismissing the guidelines as unnecessary or overzealous.

Finally, we would suggest incorporating examples of "bright patterns" in future drafts. As presently written, the guidelines include myriad examples of what not to do but no examples of what to do. As user interface design is a praxis-based field, it would be helpful to include examples of good user interface design in addition to bad user interface design. One example likely to be readily identifiable to user experience (UX) designers is the extra decision friction GitHub introduces when one deletes one's repository. To prevent accidental deletion, GitHub requires users to type the name of their repository verbatim in order to confirm the deletion. This serves the interests of users by avoiding accidental deletion even though such decision friction could easily exist as a "dark pattern." We fully recognize that the boundaries between "bright patterns" (e.g., helpful friction) and "dark patterns" (e.g., unhelpful friction) are hard to draw, and we hope any boundaries drawn would be the subject of additional comment periods.

§ 2: Specific Comments

Given our experience with the widely-acknowledged failings of privacy self-management,[7] we have focused our specific comments on areas in the guidelines that deal with such self-management. In particular we focus below on paragraphs 28, 121, 124, Annex § 4.6, and Use Case 3b.

1. **Paragraph 28: Clarify the definition of Continuous Prompting**

With regard to the definition of Continuous Prompting, we believe the following change should be made:

> "[...] by being repeatedly asked to provide additional data ~~and~~ **or** offered arguments why they should provide it."

The original wording requires both repeated prompting and offering arguments to qualify an interface as a dark pattern. However, either one of these methods on their own may be used to push users to provide more personal data than necessary. Additionally, whether designers can offer multiple independent arguments for why users should provide data in a single request to provide data is unclear. That leaves designers without crucial guidance on whether they can make multiple arguments in a single instance of asking for data. Likewise, the definition uses "repeatedly" without specifying what would qualify, opening the possibility for abuse from subjective interpretations. The guidelines could, and should, provide more detail on how "repeatedly" should be interpreted. In short, how many times count as "repeatedly?"

Furthermore, it is unclear to what extent (if any) the Continuous Prompting dark pattern applies across devices. For example, if one denies providing one's personal data to a platform on one's computer, but one is still asked for one's personal data when browsing the same platform on other devices (e.g., mobile phone, smart-TV, IoT device, voice assistant, etc.), is one being subjected to Continuous

[7] *See generally,* Daniel J. Solove, *Introduction: Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880 (2013).

Prompting? The answer should be yes. We suggest a clarification be added that prompting a user multiple times, even if done only once per device or discrete user interaction[8], could still constitute a dark pattern. This would encourage the kind of holistic thinking that prevents a myopic focus on making the user experience seamless that often underlies pernicious design practices.

2. Paragraph 121: Consolidate Settings using a Bright Line Rule

We believe that the guidelines around 4.1.2, the "Privacy Maze," should explicitly encourage platforms to consolidate their privacy settings. The guidelines could go even further by specifying standardized rules over settings that relate to user data.[9] Moves towards standardization would also improve the ability to audit platform compliance since relevant settings could be found in the same place across platforms. Improving recognizability and understanding of user data settings through standardization would especially benefit marginalized users, who are currently impacted the most by information asymmetry.[10] An easy way to set a bright line rule in this context is to suggest, possibly in paragraph 116, that all user data settings be placed at the same website location, like [example.com]/data-settings, with clear signposting for users.

3. Paragraph 124: Provide a better example of Lacking Hierarchy

While the birthday example provided demonstrates an example of a Lacking Hierarchy (inconsistent ordering) of data privacy options, it creates serious user burdens to ask users to tweak the privacy settings at such a granular level (e.g., year, month, date, time of birthday). Though admittedly some users may want different privacy settings to apply to the year, month, and date of their birthday, flooding users with granular choices, even if generally autonomy-enhancing, can be a burdensome dark pattern in itself to the extent users feel rushed or overwhelmed.[11] Users are likely, in such a situation, to stick to a default option. From the design perspective though, the example is simply unrealistic. Since it is unlikely to enhance user satisfaction, it is unlikely to be created, limiting the ability for designers to apply the lessons of Use Case 3b to interfaces they are actually likely to create. We suggest providing an example that is more relevant to what users may encounter.

4. Section 4.6: Create 4.6.4 "Misleading Necessity"

---

[8] In fact, prompting across multiple devices is a common feature of the business strategy known as omni-channel marketing. *See* Elizabeth Manser Payne, James W. Peltier & Victor A. Barger, *Omni-channel marketing, integrated marketing communications and consumer engagement: A research agenda*, 11 J. OF RESEARCH IN INT. MARKETING 185 (2017).

[9] Habib, Hana et al., "" It's a scavenger hunt": Usability of Websites' Opt-Out and Data Deletion Choices." *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020 (providing arguments for consolidated settings in a standardized location, among other best practices, based on their user studies).

[10] Alice E. Marwick, *Privacy at the margins: Introduction*, 12 INT'L J. OF COMM. 9 (2018).

[11] Examples of work that explores ways in which the burden of additional user privacy settings can be mitigated includes,

- Ravichandran, Ramprasad, et al. "Capturing social networking privacy preferences." *International symposium on privacy enhancing technologies symposium*. Springer, Berlin, Heidelberg, 2009.
- Smullen, Daniel, et al. "The Best of Both Worlds: Mitigating Trade-offs Between Accuracy and User Burden in Capturing Mobile App Privacy Preferences." *Proc. Priv. Enhancing Technol.* 2020.1 (2020): 195-215.

Without question, platforms should make transparent what data is necessary for the desired functioning of a service versus what data the platform *wants* for general processing. However, we believe this is a problem distinct from more general categories of leaving consumers in the dark. While it is a dark pattern that relies on an information asymmetry and a deliberate obfuscation on the part of platforms, it is not akin to language discontinuity, conflicting information, or ambiguous wording in that it often works by omission and relies on the assumptions of users. For example, when setting up an Instagram account, one only needs to provide a name and a phone number or email in order to acquire an account but many users are also asked for access to their contact information and may assume that is necessary for the application to function.

5. Use Case 3b: Managing One's Data Protection Settings

The issues we have with Use Case 3b stem largely from the designer's perspective. From a designer's perspective, a major question raised in this Use Case is what to do when two individuals' settings conflict. A good example of this is what to do when the "poster's" settings conflict with the "liker's" settings. Generally speaking, the more a designer has to negotiate and resolve conflict between multiple settings, the more likely platform performance will degrade and the less likely it is profit-maximizing to respect individual choices instead of applying an across-the-board rule. More importantly, how giving the user the opportunity to manage one's own data protection settings interacts with a given platform's terms of service is not addressed. Any potential conflict between company rules and the guidelines is likely to be at the forefront of designers' minds. Not addressing these conflicts may be a missed opportunity.

* * *

As social media platforms use dark patterns to increase user engagement, data collection, and financial gain, the European Data Protection Board can serve as an important source of protection for users. In § 1 of this comment, we make the following general suggestions: clarify central definitions (i.e., personal data, dark patterns, and social media), consider different guidelines for users and practitioners, and incorporate examples of non-manipulative user interfaces (i.e., "bright patterns"). § 2 includes specific suggested changes to the guideline's main text aimed at improving readability and persuasiveness. We thank the members of the EDPB for their purposeful and critical efforts. Meaningful regulation is necessary to protect consumer autonomy and welfare. We are available to assist the EDPB towards the goal of protecting user autonomy and data rights amidst dark interfaces, and we look forward to future drafts of the guidelines.

Respectfully submitted,

Jake Chanenson
Graduate Student, Department of Computer Science, University of Chicago

Marshini Chetty
Assistant Professor, Department of Computer Science, University of Chicago

Chris Crum
Law Student, University of Chicago Law School

Nick Feamster
Neubauer Professor, Department of Computer Science, University of Chicago

Justin Holiman
Law Student, University of Chicago Law School

Ambika Khanna
Graduate Student, Harris School of Public Policy, University of Chicago

Kyle MacMillan
Graduate Student, Department of Computer Science, University of Chicago

Bishwa Pandey
Graduate Student, Harris School of Public Policy, University of Chicago

Angela Peterson
Law Student, University of Chicago Law School

Zachary Rothstein
Undergraduate Student, Department of Computer Science, University of Chicago

Brennan Schaffner
Graduate Student, Department of Computer Science, University of Chicago

Lior Strahilevitz
Sidley Austin Professor of Law, University of Chicago Law School

Joshua Yasmeh
Law Student, University of Chicago Law School

Contact: ccrum@uchicago.edu ♢ bschaffner@uchicago.edu