



An Ontology of Dark Patterns Knowledge: Foundations, Definitions, and a Pathway for Shared Knowledge-Building

Colin M. Gray
comgray@iu.edu
Indiana University
Bloomington, Indiana, USA

Cristiana Teixeira Santos
c.teixeirasantos@uu.nl
Utrecht University
The Netherlands

Nataliia Bielova
nataliia.bielova@inria.fr
Inria Centre at Université Côte d'Azur
France

Thomas Mildner
mildner@uni-bremen.de
University of Bremen
Germany

ABSTRACT

Deceptive and coercive design practices are increasingly used by companies to extract profit, harvest data, and limit consumer choice. Dark patterns represent the most common contemporary amalgamation of these problematic practices, connecting designers, technologists, scholars, regulators, and legal professionals in transdisciplinary dialogue. However, a lack of universally accepted definitions across the academic, legislative, practitioner, and regulatory space has likely limited the impact that scholarship on dark patterns might have in supporting sanctions and evolved design practices. In this paper, we seek to support the development of a shared language of dark patterns, harmonizing ten existing regulatory and academic taxonomies of dark patterns and proposing a three-level ontology with standardized definitions for 64 synthesized dark pattern types across low-, meso-, and high-level patterns. We illustrate how this ontology can support translational research and regulatory action, including transdisciplinary pathways to extend our initial types through new empirical work across application and technology domains.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**; **Empirical studies in HCI**.

KEYWORDS

dark patterns, deceptive design, regulation, ontology

ACM Reference Format:
Colin M. Gray, Cristiana Teixeira Santos, Nataliia Bielova, and Thomas Mildner. 2024. An Ontology of Dark Patterns Knowledge: Foundations, Definitions, and a Pathway for Shared Knowledge-Building. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24), May 11–16, 2024, Honolulu, HI, USA.* ACM, New York, NY, USA, 22 pages. https://doi.org/10.1145/3613904.3642436

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '24, May 11–16, 2024, Honolulu, HI, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0330-0/24/05
https://doi.org/10.1145/3613904.3642436

1 INTRODUCTION

Deceptive design practices are increasingly common in digital environments, impacting digital experiences on social media [42, 51], e-commerce [40], mobile devices [29], cookie consent banners [26], and gaming [58], among others. An increasingly dominant framing of these deceptive practices is known as "dark patterns"[1]—describing instances where design choices subvert, impair, or distort the ability of a user to make autonomous and informed choices in relation to digital systems regardless of the designer's intent [11, 14, 21].

While the origins of dark patterns as a concept to describe manipulative design practices goes back over a decade to when the term was coined by practitioner and scholar Harry Brignull [6], in the past five years there has been growing momentum in the use of the term to unite scholars, regulators, and designers in transdisciplinary dialogue to identify problematic practices and find ways to prevent or discourage the use of these patterns. In 2021, Mathur and colleagues [41] published a paper at CHI beginning this work in uniting the community by outlining the general scope of the term "dark patterns," proposing common attributes, and identifying methods for identifying and characterizing dark patterns—a paper which has since supported regulatory and legal action relating to dark patterns. This momentum is also borne out on social media; according to a recent study of the historical evolution of #darkpatterns on Twitter (since renamed to "X") by Obi and colleagues [47], since 2019, conversations have included stakeholders not only from design and technology but also social scientists, lawyers, journalists, lawmakers, and members of regulatory bodies and consumer protection organizations.

Within the regulatory space, in 2022 alone, the term "dark patterns" was codified into EU law in the Digital Services Act [14], the Digital Markets Act [13], and the Data Act proposal [12], and into US law in the California CPRA [11]. Regulatory bodies such as the US Federal Trade Commission (FTC), the UK Competition and Market Authority (CMA), the EU Commission, the European Data Protection Board (EDPB) and the the Organisation for Economic

[1]We use this term to connect our efforts to prior scholarship and legal statute, recognizing that other terms such as "deceptive design" or "manipulative design" are sometimes used to describe similar tactics. While the ACM Diversity and Inclusion Council has included dark patterns on a list of potentially problematic terms, there is no other term currently in use that describes the broad remit of dark patterns practices that include deceptive, manipulative, and coercive patterns that limit user agency and are often hidden to the user.

Co-operation and Development (OECD) have released guidance on specific types of dark patterns with various levels of overlap with definitions from academic scholarship [10, 15, 16, 21, 48]. In late Summer 2023, the Department of Consumer Affairs in India also released draft guidelines regarding dark patterns [33] which were finalized in November 2023 [1]. In addition, the concept of dark patterns has been leveraged in sanctions against companies that have relied upon manipulative practices. Recent actions include a $245 million USD judgment against Fortnite, a product from Epic Games, for their use of manipulative practices to encourage the purchase of content [56] and multiple settlements by various US states against Google for their use of dark patterns to obtain location data [49, 55]. In the EU, both Data Protection Authorities (DPAs) and court decisions have forbidden certain practices related to dark patterns, including: pre-selection of choices [9]; refusing consent if it is more difficult than giving it [19, 20]; and misinforming users on the purposes of processing data and how to reject them [20, 39].

As part of this convergent discourse, HCI scholars have addressed the threat of dark patterns in a wide range of publications, proposing definitions and types of dark patterns [4, 24, 37, 40, 41]. However, the specific forms that dark patterns can take, the role of context, the ubiquity of the practices, the technologies used or application area, the comparative harms of different patterns, remedies, and the role of user education and countermeasures are still a topic of ongoing research. The consequence of this dynamic topic is of an ever-expanding list of categories and variants whose scale continues to grow.

Two large challenges face an ongoing transdisciplinary engagement with the concept of dark patterns. First, the literature has grown quickly and tends to be siloed, often lacking accurate citation provenance trails of given typologies and definitions, making it difficult to trace where new or more detailed types of patterns emerged and under which conditions. For instance, some patterns were originally coined in particular framings of user interaction such as privacy (e.g., Bösch's "immortal accounts") or rely on domain-specific characteristics (e.g., Brignull's "sneak into basket" which is strongly associated with e-commerce). Without these original contexts in mind, it can be difficult to understand how the use of a pattern and its associated definition can be productively (or problematically) applied to a new domain. In parallel, the space that dark patterns scholars have sought to cover is also vast, with important research occurring in specific domains (e.g., games, e-commerce, privacy and data protection) and across different technologies and modalities (e.g., mobile, desktop, conversational user interfaces (CUIs), AR/VR), as shown in a recent systematic review of dark patterns literature [22]. This diversity of research has led some scholars to propose fragmentary, domain-specific typologies without necessarily finding commonalities across domains—resulting in extra work and often needlessly strengthening scholarly siloes. Second, regulators and policy makers have been interested in the scholarly conversation regarding dark patterns, but have in some cases created wholly new domain-related terminology to describe types already known in the academic literature (e.g., the EDPB social media guidelines [16, 17], which included many previously known dark patterns that were described by wholly new names, severing connections to other relevant literature) in their pursuit of providing legal guidance on emergent issues relating to dark patterns (e.g., [17]. In other cases, regulators and policymakers have inconsistently cited academic sources (e.g., [17, 21]) making connections across the regulatory, legal, and academic spaces fraught—and making academic and practitioner work that connects these domains difficult to broker.

We seek to support these challenges and ongoing conversations by building the foundation for a common ontology of dark patterns. This effort is directly motivated by multiple years of engagement by the research team—including discussions with participants at numerous international conference presentations, workshops, and symposia, alongside interactions with regulators, legal scholars, and engagement as expert witnesses on legal cases relating to dark patterns. Through these encounters, we have confronted both the challenges of conducting work in an emergent space where there is broad consensus on the key components of dark patterns but not necessarily a shared language (as extensively described by Mathur and colleagues [41]) and the promise of synergies with other interdisciplinary partners when this shared language is realized. In particular, there has been broad interest in using a consolidated set of terms to describe the types of dark patterns, their presence, and their impacts—connecting the design of digital systems, social scientists that study the implications of these systems on users, and regulators that seek to rein in unfair, deceptive, or coercive business practices—regardless of the domain or context in which these practices occur.

By taking the first steps towards building an ontology, we seek to create a shareable, extendable, and reusable knowledge representation of dark patterns which is hosted at **https://ontology.darkpatternsresearchandimpact.com**. This groundwork for an ontology is both domain and application agnostic though it has potential utility in domain or context-specific instances as well. For instance, the *Bad Defaults* dark pattern is often embedded in settings menus, pre-set so that users share personal information on social media platforms or accepting to receive advertising content on online shopping sites unknowingly. Such context-specific instances are enabled through *Interface Interference*—a domain-agnostic strategy used to manipulate interfaces, privileging certain actions and, thus, limiting discoverability of alternatives. As noted by Fonseca [18], ontologies can be useful in supporting social science research by "creating better conceptual schemas and applications." We build upon this argument from Fonseca, arguing that our ontology also supports alignment across social science researchers, legal scholars, regulators, and designers—supporting these stakeholders with a shared vocabulary which they can use to discuss existing and emergent concerns relating to dark patterns across a variety of domains and contexts. To create this preliminary ontology, we build upon ten contemporary taxonomies of dark patterns from both the academic and regulatory literature, and thereafter we identify three levels of hierarchy for pattern types. Hence we harmonize concepts across these taxonomies to provide a consistent and consolidated, shared, and reusable dark patterns ontology for future research, regulatory action, and sanctions.

We make four contributions in this paper. First, we introduce the hierarchical concepts of low-level, meso-level, and high-level dark patterns to the literature, disambiguating UI-level patterns that may lead to opportunities for detection (low-level) and strategies

that may be targeted by policy and legislation (meso- and high-level). Second, by analysing the provenance of dark patterns from academic and regulatory sources, we identify when patterns first emerged and how naming has evolved over time and across sources. Third, we describe a common definition syntax, set of definitions, and hierarchy of dark patterns that aligns disparate terminology from scholars and regulators. Fourth, we demonstrate how the ontology can be strengthened and extended through additional empirical work, and how the ontology can effectively be utilized by practitioners, scholars, regulators, and legal professionals to support transdisciplinary action.

2 MOTIVATION & BACKGROUND

Since the initial set of a dozen types of dark patterns proposed by Brignull in the 2010s, research has focused on related issues from multiple angles including, but not limited to, e-commerce, games, social media, and IoT [22]. While this scholarship contributes significant insights to the discourse, we noticed varying approaches to adopt existing descriptions, defining novel scenarios in which users are harmed. Meanwhile, the specification of individual typologies creates a certain ambiguity within the overall discourse on the matter. In developing this ontology, we confront numerous timely issues relating to the description of dark patterns, the study of dark patterns and their harms through empirical work, and the leveraging of this scholarship to support legal and regulatory action.

Dark patterns are known to be ubiquitous; however, most pattern types have been explored in relatively narrow contexts, cultures, or domains with more scholarship needed to fully define causal links, harms, and impacted populations [22]. The HCI community has been engaged and interested in impacting society and the future of technology practices relating to dark patterns [23, 27, 38]—and indeed, HCI scholars have been central in the study of dark patterns, revealing insights relating to the harm and severity of dark patterns that then support enforcement action and regulation. However, we currently lack a shared landscape of definitions, types, and language to unify the study of dark patterns. Without this shared landscape, research has become (and will continuously be) fragmented by domain, context, and technology type—which if not addressed, may lead to duplicated effort by scholars working on similar issues in different domains, and additionally may hamper regulatory enforcement due to lack of precision and shared language regarding precisely what dark patterns are used and with what effect. Such lack of a shared ontological framework may also restrict traceability and searchability of dark patterns.

Our work unifies practitioner, scholarly, and regulatory efforts that describe the range of dark patterns, leading to a shared vocabulary and ontology that allows for coordination of efforts across diverse contexts (e.g., technologies, specific functionality, areas of technology use) and stakeholders (e.g., regulators, legal scholars, social scientists, practitioners). This ontology will support not only the advancement of scholarship, but also translational and transdisciplinary efforts that connect scholarship to legal sanctions and regulatory frameworks. For instance, there are now high-level prohibitions of dark patterns by regulatory authorities and legal statute; however, the specific low-level practices that should be deemed illegal under these prohibitions are not yet detailed in enforcement action or case law. This paper connects these different strands of work by harmonizing regulatory and academic work into a single ontology, enabling future scholars and practitioners from all disciplines to utilise our structures and definitions to support their work.

3 METHODOLOGY

We used a qualitative content analysis approach [31] to identify and characterize elements of existing dark patterns taxonomies using the method described in Figure 1[2].

As a research team, we leveraged our collective experiences in human-computer interaction, design, computer science, law, and regulation. Specifically, our team included established dark patterns scholars, including one with a focus on human-computer interaction and design (Gray and Mildner), one with a focus on computer science and web measurement and experience in regulation (Santos), and one with a background in computer science and data protection law (Bielova). Across these perspectives, in accordance with previous scholarship, we sought to characterize dark patterns in a transdisciplinary way, drawing on multiple disciplinary perspectives that provide differing views on the origins and types of dark patterns [26]. However, these backgrounds also introduce gaps, tensions, and opportunities that relate to the unique experience and academic training of each author. To account for this difference in perspective, each dark pattern type was initially reviewed by each author independently before engaging in conversation amongst the researchers that led to the final agreement on the harmonized type and definition.

3.1 Data Collection

We collected dark patterns taxonomies in Fall 2022 from a total of 10 sources, including:

(1) A set of patterns shared on https://darkpatterns.org since 2010 by Harry Brignull[3].
(2) Scholarly academic sources that were highly cited[4], present a distinct and comprehensive taxonomy, and have had one or more component dark patterns types included in a regulatory reports (with or without citation) [4, 24, 37, 40].
(3) Public reports from stakeholders and regulators in the EU, UK, and USA that include a dark patterns taxonomy [10, 15, 16, 21, 48].

The selection of these sources encompass, at the time of our data collection in Fall 2022: i) the "classic" set of patterns shared on darkpatterns.org for over a decade by Brignull; ii) the most commonly cited taxonomies in the research literature (which were also referenced in regulatory taxonomies in a direct or indirect way, likely due to their prominence), and iii) the most comprehensive set

[2]This work builds upon an extends a previous draft version of this ontology published at CHI 2024 as a late-breaking work. [25]
[3]This collection of dark patterns was moved to https://www.deceptive.design in 2022, but the 12 patterns we drew on have been stable since 2018 when the final pattern, "confirmshaming," was added. In 2023, this website was updated to include additional pattern types, resulting in a modified collection of 16 types.
[4]As of December 2023, Google Scholar reports citations of these sources as follows: Gray et al. [24] (657); Mathur et al. [40] (454); Luguri and Strahilevitz [37] (260); and Bösch et al. [4] (259).

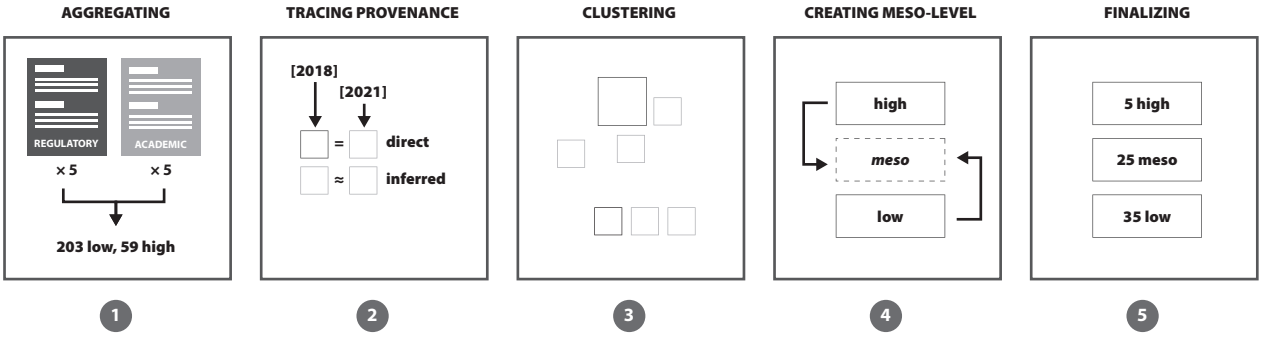


Figure 1: Our method for creating the ontology, mapping to the steps in Section 3.2.1.

of regulatory literature from countries that had produced reports specific to dark patterns at this time. Taken as a set, these academic, practitioner, and regulatory taxonomies provide a strong foundation for our ontology, setting the stage for inclusion of other domain- or context-specific taxonomies in the future (which we outline in Section 6.

3.2 Data Analysis

Once we gathered the set of taxonomies, we began our analysis by identifying the constitutive components of each taxonomy without considering overlaps across sources through a bottom-up approach.

Quantification of dark pattern types Across the ten taxonomies from academic and regulatory sources collected in Fall 2022, we identified 186 low-level and 59 high-level patterns (a total of 245 patterns).

After our initial analysis, the patterns used on Brignull's site (https://www.deceptive.design) were substantially updated in the Summer 2023, and we collected the additional set of patterns for that source—resulting in 11 total sources. Also, the EDPB regulatory report was made final in February 2023, and we used its final taxonomy in this paper after completing our initial mapping in the Fall 2022 based on the draft report taxonomy. Based on the updates to the EDPB guidelines and Brignull's site in the Spring and Summer 2023, the total number of patterns we analyzed included 203 low-level (adding 1 new pattern from the revised EDPB guidelines and 16 patterns from the updated Brignull site) and 59 high-level patterns—**a total of 262 patterns** (see Tables 2 and 3 in the supplemental material). All taxonomy elements are included in supplemental material for other scholars to build upon.

Rationale underlying the high number of dark pattern types This large number of discrete elements is perhaps unsurprising, since each typology author has used a different point of focus and categorization based on the sector they sought to describe or support. For instance, Mathur et al. [41] and the CMA [10] focus on e-commerce; the EDPB focuses on data protection practices within social media platform interfaces [16], and the FTC [21] and EU Commission [15] focus on guidance specific to their jurisdictions and underlying legal authority. The types themselves also evolved in one case due to input from the practitioner and regulatory community, which is the case of the EDPB naming of patterns changed slightly from the 2022 draft report to the final 2023 report, with one high-level strategy "hindering" changing to "obstructing" to bring it into better alignment with academic taxonomies.

3.2.1 Creating the Ontology Framework. We used the following procedure to carefully identify existing taxonomy components, their source, relationships and similarities between components across taxonomies, visualized in Figure 1:

(1) **Aggregating existing patterns.** We first listed all high- and low-level patterns verbatim in the structure originally indicated in the textual source. *High-level patterns* include any instances where the pattern is denoted as a category, strategy, goal, intention, or other parent in a parent-child relationship. *Low-level patterns* indicate specific patterns that are included as a child in a parent-child relationship, or are otherwise undifferentiated in hierarchy (e.g., Brignull's patterns).
(2) **Identifying provenance through direct citations and inference.** Based on citations provided in the source-document, we indicated any instances where patterns were *directly cited* or otherwise duplicated from previous sources. Because many patterns were uncited—particularly in regulatory reports—we also relied upon citations elsewhere in the document or explicit use of existing pattern vocabulary and definitions from previously published sources, which we indicate as *inferential*. We used these direct and inferential citation patterns to identify where patterns were first introduced, even if they appeared alongside other patterns that had been published previously. This allowed us to map the historical progression of high- and low-level types over time.
(3) **Clustering similar patterns.** We grouped patterns that appeared either to be identical or similar (in a *is-a* or *equivalent-to* relationship) on Miro (see Figure 2), using definitions to identify affinities among patterns that did not have identical names. This portion of the analysis was the most extensive, including in depth conversations between an HCI and legal scholars and a careful reading of the definitions as they might be understood by designers and lawyers. We tried out numerous different groupings based on what we understood

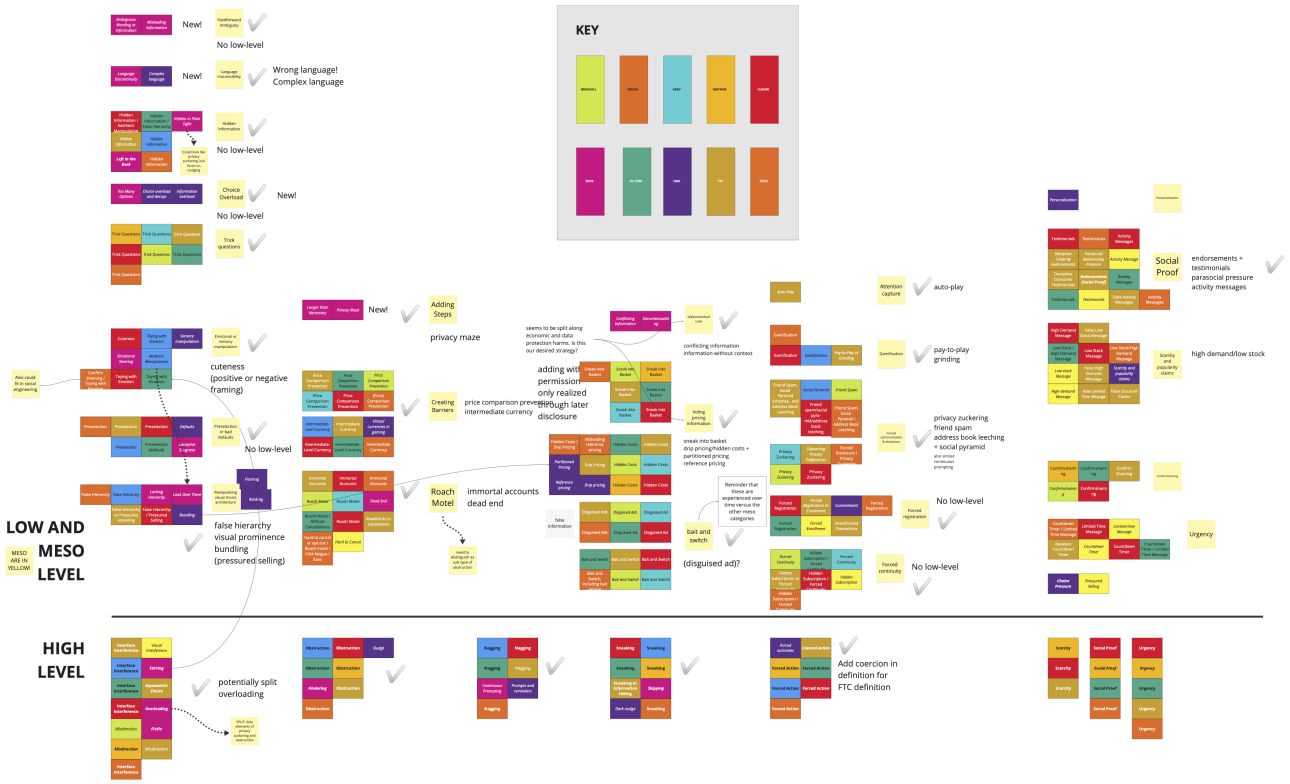


Figure 2: A screenshot of our Miro workspace where we organized and clustered elements of the ten source taxonomies. Columns indicate an entire structure of meso- and low-level patterns underneath a high-level pattern and yellow Post-It notes indicate draft meso-level patterns. The elements are color-coded based on which taxonomy they came from. A full version of this workspace is included as a supplemental material.

to be the main focus of each pattern and then sought to characterize what level of pattern each represented.

(4) **Creating meso-level patterns.** From the findings of this visually-organized analysis procedure, we recognized that there were not only low- and high-level patterns present, but also a *"meso"* level of pattern knowledge. By recognizing similarities among low-level patterns, we introduced *meso-level patterns* into our analysis, identifying these patterns by using the names or elements of existing taxonomies where possible, or coining new names to characterize the low-level patterns we grouped together. If the pattern cluster was specific to low-level UI concerns, we sought to identify a meso-level pattern name that was more abstract and could contain the low-level pattern. If the pattern represented a meso-level abstraction, we did not seek to identify specific low-level instantiations—instead leaving that task for future scholarship efforts in domain- and technology-specific areas.

(5) **Finalizing the ontology.** Across these three levels of hierarchy, we grouped 233 of the 245 taxonomy elements[5]. After evaluating the changes to the EDPB guideline taxonomy and updated Brignull taxonomy in Spring and Summer 2023, we updated our mapping of 262 patterns, which resulted in no additional novel pattern types. The final ontology includes 5 high-level patterns, 25 meso-level patterns, and 35 low-level patterns—**a total of 65 patterns**.

3.2.2 Harmonizing Definitions of Dark Patterns Types. Building on this ontology framework, we then proceeded to create a definitional syntax across the three levels of the ontology and then created definitions for each final pattern using the following approach:

(1) **Creating definition syntax.** We evaluated the range of approaches to definitions in the existing taxonomies.

[5]Four ungrouped elements were from the CMA report [10] in Fall 2022 and described generic elements of digital systems which were not explicitly framed as deceptive or manipulative: Choice Structure, Choice Information, Feedback, and Messengers. All eight high-level patterns from Bösch [4] were also excluded since they were not reiterated in any downstream literature.

- *Short vs long definitions.* Some definitions were very short (e.g., the EU Commission's definition for *forced registration*: "Consumer tricked into thinking registration is necessary") while other definitions were more elaborate (e.g., the FTC's definition for *baseless countdown timer*: "Creating pressure to buy immediately by showing a fake countdown clock that just goes away or resets when it times out. Example: 'Offer ends in 00:59:48"'; the EU Data Protection Board's definition for *longer than necessary*: "When users try to activate a control related to data protection, the user journey is made in a way that requires more steps from users, than the number of steps necessary for the activation of data invasive options. This is likely to discourage them from activating such control.").
- *Description of the definitions.* Most definitions were based in a description of user interaction with a system, like the examples above; however, Brignull's 2018 definitions were written in first-person language demonstrating how a user would experience a dark pattern (e.g., the definition for *roach motel*: "You get into a situation very easily, but then you find it is hard to get out of it (e.g. a premium subscription)."). Interestingly, Brignull's 2023 language appears to model other taxonomies with all definitions beginning with "The user struggles...," "The user expects..." or similar structures.
- *Definition structure and syntax.* We used an iterative process where two authors independently and collaboratively tested different definition structures. Based on these efforts and through discussion, we finalized sample definition structures and syntax that captured the relevant type of knowledge (e.g., strategy, angle of attack, means of execution). For instance, all high-level patterns included the interplay of an undesired action and a limitation of their decision-making or free choice. Meso-level patterns addressed a mismatch in users' expectations of a system and the relevant impact. Low-level patterns identified how they manifest their parent high- and meso-level pattern in relation to one or more elements of the UI and a mismatch of expectation and resulting effect on the user experience.

(2) **Creating and member-checking high- and meso-level pattern definitions.** We then drafted definitions for all high- and meso-level patterns, iterating on the structure until we found a syntax that appeared to address all critical elements of the existing definitions and allow us to clearly indicate how the pattern subverted user autonomy and manifest as deceptive or coercive. We began with definitions at these levels since low-level patterns were already grounded in specific UI examples, and thus more effort was needed to identify what components a definition at a higher level of abstraction should include. To support member-checking, our set of 30 definitions and the draft definition structures were then shared via a Google Doc with members of a large Slack community focused on research and enforcement action relating to dark patterns. This Slack channel was initiated in 2021 to grow and foster a community of dark pattern researchers after a successful CHI workshop on the topic [38]. The community has since grown into a transdisciplinary network with over 100 participants around the world, including dark patterns scholars, practitioners, legal experts, regulators, and representatives of non-profits seeking to combat deceptive design practices. We asked this community for feedback on the utility of the definitions, the completeness of the definition structures, and the ability of these definitions to leave as open-ended the many different low-level manifestations of dark patterns. More than two dozen community members viewed the draft materials (evidenced through comments or reactions), and over ten gave us feedback. Most interactions were quite short (particularly on Slack), while others involved threaded messages with replies to clarify meaning. For instance, there were discussions of how central "deception" should be in the definitions, requests for more information on how the different levels of definitions functioned, and specific words that can or should be used to indicate curtailing of user autonomy. Regardless of the form of interaction, the feedback was overwhelmingly positive, with respondents mentioning the utility of the patterns and definitions in supporting future work and validating our approach to focus on mechanisms that support the power of dark patterns rather than overly focusing on intent. This positive feedback mirrored what we have heard from scholars and regulators when presenting draft versions of the ontology in various international symposia, workshops, and conferences for almost a year.
While more formal evaluation of specific definitions in particular contexts (e.g., design, regulation, research) will be useful, the community members who stand to benefit the most from the ontology have ensured that the definitions have face validity. Since this initial evaluation in Summer 2023, the ontology has also been used to support an emerging collection of "fair patterns" as well, wherein meso- and high-level dark patterns types from our ontology are linked to specific countermeasures that could be considered by designers and regulators.[6]

(3) **Finalizing low-level pattern definitions.** After mapping out the initial 30 definitions, we created definitions for the 34 low-level patterns that were grounded in the specifics of the UI execution. These patterns were easier to write since many taxonomy definitions (in particular those from Brignull [7], Gray [24], and the FTC [21]) included richer detail for patterns that pointed towards a real-world implementation. As a research team, we read and edited the definitions until we were satisfied with their level of consistency and relationships to the higher-level categories in which they belonged. All definitions are included in the appendix of this paper and supplemental materials to support future work.

4 MAPPING THE EVOLUTION OF DARK PATTERNS

Pattern names have largely stabilized in the past five years, including high-level pattern types (e.g., nagging, obstruction, sneaking, interface interference, forced action) and low-level patterns (including Brignull's [6, 7] and those introduced by Gray et al. [24] and

[6] https://fairpatterns.com/what-are-dark-patterns/

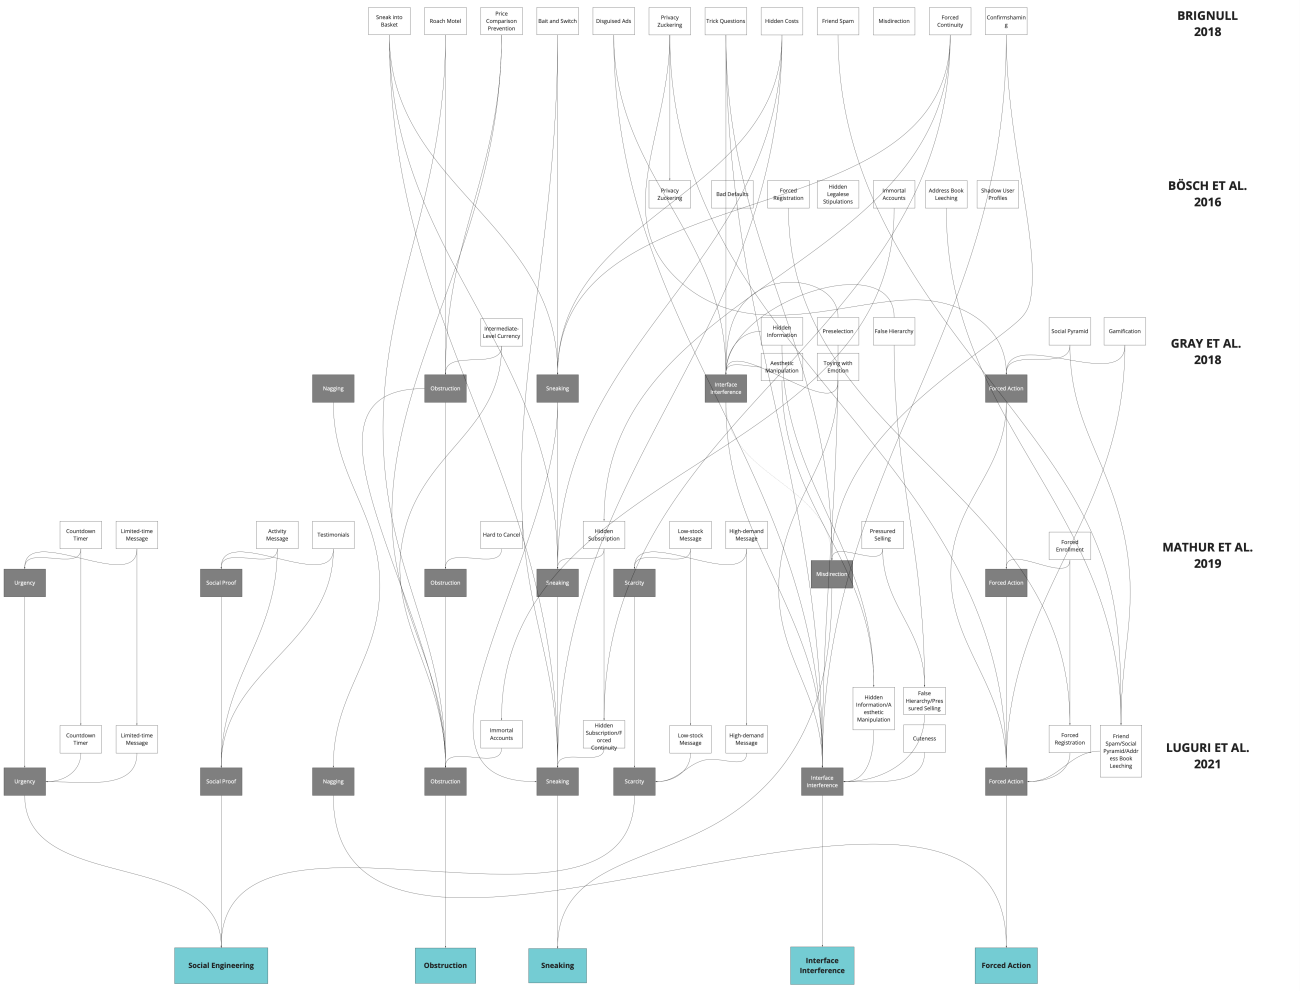


Figure 3: A visual mapping of the evolution of dark patterns in the academic taxonomies we analyzed from 2018-2021. Each row includes elements of the related taxonomy by year and source, and connecting lines indicate relationships between or reiterations of different patterns over time. Pattern names in gray boxes are high-level patterns, pattern names in white boxes are low-level patterns or otherwise lack hierarchy, and pattern names at the bottom are the final high-level patterns we adopt in our ontology. A full version of this mapping is included as a supplemental material.

Mathur et al. [40]). A mapping of these patterns over time across the academic and practitioner sources we considered is included in Figure 3.

High-level patterns were most likely to co-occur across multiple sources. For instance, Gray et al.'s [24] original five high-level "dark pattern strategies" were found across multiple other sources, even if they were not consistently cited: nagging [15, 37], obstruction [15, 21, 37, 40], sneaking [15, 21, 37, 40], interface interference [15, 21, 37], and forced action [15, 21, 37, 40] (FTC uses "coerced action" instead). As shown in Figure 3, virtually all of the high level patterns proposed by Gray et al. in 2018 were carried forward in other academic taxonomies. In Brignull's 2023 changes to https://www.deceptive.design, multiple high-level strategies from Gray et al.'s [24] taxonomy were added to the website (nagging, obstruction, sneaking, forced action, visual interference)—however, these changes were not cited and Brignull continued his practice of not providing direct citations or hierarchical structure to his patterns. After their introduction in Mathur et al. [40], newly introduced categories relating to social psychology or behavioral economics also became common: urgency [15, 21, 37], scarcity [21, 37], and social proof [15, 21, 37] (the FTC bundles "Endorsements" with "social proof"). We have grouped these types together as part of a sixth high-level pattern of "social engineering."

Domain or context-specific patterns. The most volatility has occurred in relation to *domain-* or *context-specific patterns*. These include expansions of Mathur et al.'s [40] high-level patterns of "social proof" and "scarcity," which have since been reiterated by the EU Commission [15] and OECD [48] and extended by the CMA [10] and FTC [21] taxonomies. In addition, the EDPB guidance on dark patterns in social media [16] included a wholly new set of 6 high-level and 15 low-level patterns, although the majority of these could be inferred as similar to already existing patterns proposed in the academic literature. Importantly, though, the EDPB taxonomy included multiple patterns which we found to be new low-level or meso-level additions, including "privacy maze," "dead end," "conflicting information," "information without context" (which we renamed from the EDPB pattern "decontextualizing"), and "visual prominence" (which we renamed from the EDPB pattern "look over there"). Similarly, the CMA taxonomy focused on choice architecture as a guiding structure with three categories focused on choice "structure," "information," and "pressure." This taxonomy structure also yielded new patterns, including "bundling," "complex language," and "personalization."

Our analysis demonstrates the value in classifying or generating context-specific patterns that illuminate gaps in current taxonomies, and also the benefit of mapping these patterns within larger ontologies to identify abstractions of patterns that may apply across many domains, contexts, and legal fields. Our final ontology mapping is included in Figures 4 and 5 and can also be found in the supplementary materials.

5 CREATING A DEFINITIONAL STRUCTURE BY ONTOLOGY LEVEL

As described in Section 3.2.1, our ontology includes three different levels of hierarchy:

- **High-level patterns** are the most abstracted form of knowledge, including general *strategies* that characterize the inclusion of manipulative, coercive, or deceptive elements that might limit user autonomy and decision making. These patterns are context-agnostic and can be employed through a range of modalities and technologies (e.g., desktop, mobile, VUIs, VR/AR) and application types (e.g., e-commerce, gaming, social media).
- **Meso-level patterns** bridge high- and low-level forms of knowledge and describe an *angle of attack* or specific approach to limiting, impairing, or undermining the ability of the user to make autonomous and informed decisions or choices. These patterns are content-agnostic and may be interpreted in a contextually-appropriate way based on the specific context of use or application type.
- **Low-level patterns** are the most situated and contextually dependent form of knowledge, including specific *means of execution* that limits or undermines user autonomy and decision making, is described in visual and/or temporal form(s), and is likely to be detectable through algorithmic, manual, or other technical means.

To create a definitional structure for each level, we first used a subset of approximately ten dark patterns types and definitions in order to "play-test" a combined and unified definition for dark patterns types at multiple levels of granularity (i.e., high, meso, low). Through this process, we considered not only the level of abstraction inherent in dark patterns at differing levels, but also the interaction between: the user's expectations of what should or would be likely to occur (i.e., manipulation of the gulf of execution); the user's identification that something had occurred that they did not wish to happen (i.e., manipulation of the gulf of evaluation); and the mechanisms used to inform or execute manipulation in either of these prior elements. We also considered cases where the deception or manipulation was likely to be hidden to the user (e.g., cases of sneaking, obstruction, or interface interference) as well as cases where deception or coercion was overt and known to the user (e.g., forced action). Based on this iterative generation of a definitional structure, we created a standardized syntax for each dark pattern level, described below. All 65 final definitions are included as a supplemental material.

5.1 High-Level Patterns

> **{HIGH-LEVEL DARK PATTERN}** is a strategy which **{UNDESIRED ACTION}** that [optionally, if known to users, would] **{DISTORT/SUBVERT/IMPEDE/OTHERWISE LIMIT USERS' AUTONOMY, DECISION-MAKING, OR FREE CHOICE}**.

Across our 5 high-level pattern definitions, we considered *undesired actions* such as: hiding, disguising, delaying, redirecting, repeating, impeding, privileging, or requiring actions. We also considered a range of *mechanisms* that could be used to limit users' autonomy, decision-making, or free choice such as: foregrounding unrelated tasks, dissuading a user from taking an action, confusing the user, limiting discoverability of action possibilities, causing a user to unintentionally take an action they would likely object to, or forcing a user to take an action they would not otherwise take. Most of these definitions placed a focus on mechanisms which were primarily hidden, resulting in the user being deceived, such as: "*Interface Interference* is a strategy which privileges specific actions over others through manipulation of the user interface, thereby confusing the user or limiting discoverability of relevant action possibilities." However, the definition for *Forced Action* was focused more on the coercive nature of the interaction which may involve users' awareness they are being manipulated: "*Forced Action* is a strategy which requires users to perform an additional and/or tangential action or information to access (or continue to access) specific functionality, preventing them from continuing their interaction with a system without performing that action."

5.2 Meso-Level Patterns

> **{MESO-LEVEL DARK PATTERN}** subverts the user's expectation that **{EXPECTATION}**, instead producing or informing **{DIFFERENT EFFECT ON USER}**.

Across our 24 meso-level pattern definitions, we considered a range of *user expectations* such as: presence of relevant and timely information, match between user goal and action, completeness and truthfulness of information provided, and the ability to change one's mind and reverse a decision. We also considered a range of *potential*

High-Level Pattern	Meso-Level Pattern	Low-Level Pattern
Obstruction D: Gr Lu Ma Br23 EUCOM FTC OECD I: EDPB CMA	Roach Motel (D: Br Gr Lu EUCOM I: Br23 Ma FTC OECD)	Immortal Accounts (D: Bö Lu FTC OECD)
		Dead End (D: EDPB)
	Creating Barriers	Price Comparison Prevention (D: Br Gr Lu FTC EUCOM OECD; I: Br23)
		Intermediate Currency (D: Gr Lu FTC EUCOM OECD; I: CMA)
	Adding Steps (I: EDPB)	Privacy Maze (D: EDPB)
Sneaking D: Gr Lu Ma EUCOM OECD I: EDPB CMA FTC	Bait and Switch (D: Br Gr Lu FTC EUCOM I: OECD)	Disguised Ad (D: Br Gr Lu FTC EUCOM OECD; I: Br23)
	Hiding Information	Sneak into Basket (D: Br Gr Ma Lu FTC EUCOM OECD)
		Drip Pricing, Hidden Costs, or Partitioned Pricing (D: Br Br23 Gr Ma Lu CMA FTC EUCOM OECD)
		Reference Pricing (D: CMA OECD)
	(De)contextualizing Cues	Conflicting Information (D: EDPB)
		Information without Context (I: EDPB)
Interface Interference D: Gr Lu EUCOM FTC OECD I: Br Ma EDPB FTC	*Manipulating Choice Architecture* (I: CMA)	False Hierarchy (D: Gr OECD I: Lu EDPB FTC)
		Visual Prominence (I: EDPB)
		Bundling (D: CMA)
		Pressured Selling (D: Ma; I: Lu FTC)
	Bad Defaults (D: Bö; I: CMA EUCOM)	–
	Emotional or Sensory Manipulation (I: Gr Lu EUCOM OECD)	Cuteness (D: Lu)
		Positive or Negative Framing (I: Gr Lu EDPB)
	Trick Questions (D: Br Gr Ma Lu FTC EUCOM OECD; I: Br23)	–
	Choice Overload (I: EDPB CMA)	–
	Hidden Information (D: Gr FTC OECD; I: Lu Bö EDPB EUCOM)	–
	Language Inaccessibility	*Wrong Language* (I: EDPB)
		Complex Language (D: CMA)
	Feedforward Ambiguity (I: EDPB)	–

Figure 4: Our ontology of dark patterns organized by level of pattern. "D" indicates a *direct* use of the pattern language in the original source(s) and "I" indicates an *inferred* similarity between different terminology used across two or more pattern types. Sources are indicated by abbreviation and are colored cyan if they are regulatory reports or magenta if they are academic or practitioner sources. "Br" indicates his 2018 patterns and "Br23" indicates his 2023 patterns. *Italized pattern names* indicate new pattern types introduced in this paper while all other text relies upon the sources indicated. Underlined sources indicate the earliest mention of that pattern or patterns in the sources we analyzed. A full description of the inferred pattern names is included in supplemental material to support future work.

High-Level Pattern	Meso-Level Pattern	Low-Level Pattern
Forced Action D: Gr Lu Ma EUCOM OECD I: CMA FTC	Nagging (D: Gr Lu Br23 EUCOM FTC OECD; I: EDPB CMA)	–
	Forced Continuity (D: Br Gr I: Lu Ma Br23 FTC EUCOM OECD)	–
	Forced Registration (D: Bö Lu FTC EUCOM OECD; I: Bö Ma CMA FTC)	–
	Forced Communication or Disclosure	Privacy Zuckering (D: Br Bö Gr Lu; I: FTC OECD)
		Friend Spam (D: Br; I: Lu FTC OECD)
		Address Book Leeching (D: Bö; I: Lu FTC OECD)
		Social Pyramid (D: Gr; I: Lu FTC OECD)
	Gamification (D: Gr Lu OECD)	Pay-to-Play (D: FTC)
		Grinding (D: FTC)
	Attention Capture	Auto-Play (D: FTC)
Social Engineering	Scarcity and Popularity Claims (D: CMA; I: Ma Lu Br23 FTC)	High Demand (D: Ma Lu FTC EUCOM OECD)
		Low Stock (D: Ma Lu FTC EUCOM OECD)
	Social Proof (D: Ma Lu EUCOM OECD; I: Br23)	Endorsements and Testimonials (D: Ma Lu FTC EUCOM OECD)
		Parasocial Pressure (I: FTC)
	Urgency (D: Ma Lu FTC EUCOM OECD; I: Br23)	Activity Messages (D: Ma Lu FTC EUCOM OECD)
		Countdown Timer (D: Ma Lu FTC; I: EUCOM OECD)
		Limited Time Message (D: Ma Lu FTC; I: EUCOM OECD)
	Shaming	Confirmshaming (D: Br Ma Lu Br23 FTC EUCOM; I: OECD)
	Personalization (D: CMA)	–

Figure 5: Ontology of dark patterns organized by level of pattern, continued.

negative effects on the user, such as: unexpected or unanticipated outcomes, confusion or pressure, being prevented from locating relevant information, or making a different choice than they would otherwise make. Meso-level definitions as a set touched on many different aspects of the user experience, with some pointing more towards static moments in the user journey and others describing temporal effects that might be realized over a longer portion of the user journey. For instance, these two patterns represent instances where the focus was primarily on static UI elements or a particular moment of interaction:

- "*Manipulating Choice Architecture* subverts the user's expectation that the options presented will support their desired goal, instead including an order or structure of options that makes another outcome more likely."
- "*Scarcity or Popularity Claims* subverts the user's expectation that information provided about a product's availability or desirability is accurate , instead pressuring the user to purchase a product without additional reflection or verification."

In contrast, other patterns represented instances where the full effect of the pattern was felt over time and might involve multiple interactions with a system that accumulate to achieve the overall effect:

- "*Roach Motel* subverts the user's expectation that an action will be as easy to reverse as it is to make, instead creating a situation that is easy to get into, but difficult to get out of."
- "*Hiding Information* subverts the user's expectation that all relevant information to make an informed choice will be available to them, instead hiding information or delaying

the disclosure of information until later in the user journey that may have led to them making another choice."

5.3 Low-Level Patterns

{LOW-LEVEL DARK PATTERN} uses **{RELATED HIGH- AND MESO-LEVEL DARK PATTERN}** to **{ELEMENT OF UI ALTERED}**. As a result, **{INCORRECT USER EXPECTATION}** leads to **{UNDESIRED EFFECT ON USER}**.

Across our 35 low-level definitions, we considered a range of *means of execution* in the UI or user experience, such as: provision of information that is conflicting, prohibiting certain kinds of interactions, adding items without a user's knowledge, providing incomplete or misleading information, distracting a user through extraneous cues, or using social or other extrinsic pressure to steer user's decisions. These means of execution were supported by a wide range of *incorrect user expectations and related undesired effects*, including: preventing a user from making an informed choice about their privacy or purchase of a product, disclosing incomplete or misleading information that leads to choices the user would not otherwise make, or distracting a user and thus preventing them from discovering information that would be relevant to their decision. Low-level patterns all exploit the user experience in direct ways, but address different aspects of the experience:

- Focus on specific user interactions that are limited (e.g., "*Price Comparison Prevention* Creates Barriers and uses Obstruction by excluding relevant information, limiting the ability of a user to copy/paste, or otherwise inhibiting a user from comparing prices across two or more vendors. As a result, the user cannot make an informed decision about where to buy a product or service.")
- Focus on a coordinated set of user interactions that produce the desired effect (e.g., "*Privacy Mazes* Add Steps and use Obstruction to require a user to navigate through many pages a result, the user is prevented from easily discovering relevant information or action possibilities, leaving them unable to make informed decisions regarding their privacy.")
- Focus on discrete UI elements (e.g., "*False Hierarchy* Manipulates the Choice Architecture, using Interface Interference to give one or more options visual or interactive prominence over others, particularly where items should be in parallel rather than hierarchical. As a result, the user may misunderstand or be unable to accurately compare their options, making a selection based on a false or incomplete choice architecture.")
- Focus on user comprehension of the interface (e.g., "*Wrong Language* leverages Language Accessibility, using Interface Interference to provide important information in a different language than the official language of the country where users live. As a result, the user will not have access to relevant information about their interaction with the system and their ability to choose, leading to uninformed decisions.")

6 EXTENDING THE ONTOLOGY BASED ON CURRENT AND FUTURE SCHOLARSHIP

Dark patterns researchers have addressed the impact of manipulative, deceptive, and coercive design in a range of technological domains. While these efforts are important in protecting online users and identifying areas for regulatory or legal impact, the novelty and breadth of this work potentially hinders an exhaustive mapping of dark patterns onto our ontology. Building on our proposed ontology, we identify pathways for many stakeholders to contribute to the growth of ontology elements—both through the addition of new patterns and strengthening contextual or domain-specific examples of existing patterns. This extension can help not only to anchor instances of patterns from future studies in existing literature, but also to enable the scholarly community to extend or further characterize these pattern types. The ontology's stratification allows anyone to extend the current framework by following the structure and syntax given for each high, meso, and low level dark pattern type.

To perform this mapping and extension exercise, we sought to identify existing alignment between proposed dark patterns and the ontology. To this end, we consider how a source might offer new perspectives for existing or examples of novel dark patterns. The method we used to extend the ontology involves three steps:

(1) We analyzed the dark pattern definition included by the author and, if provided, considered any cited relationships to other dark patterns and related terminologies.
(2) We then aligned the author's definition with the syntax of the high, meso, and low levels, placing the dark pattern at the most logical level of abstraction.
(3) Finally, we considered how the addition of the type informs a revision of the ontology. A type could reiterate an existing type in the ontology (leaving the core ontology unchanged), extend an existing type in the ontology (providing rationale for a more expansive definition of an existing type), or identify the presence of a wholly new type (adding a type to the core ontology).

This section demonstrates how we envision for the community to extend the ontology by drawing examples from three contemporary studies defining dark patterns from domain and context-specific areas, underlining the decision behind selecting these relevant works. These examples extend across multiple emergent areas of the state-of-the-art in dark patterns literature, encompassing some of the first examples of studies addressing: dark patterns in Japanese apps (Section 6.1), dark patterns experienced across multiple modalities (Section 6.2), and dark patterns experienced on prominent social media apps (Section 6.3). We also show how the ontology can be extended to map legislation and case law relating to dark patterns. Table 1 summarizes how three different sources were compared to our ontology through this method, demonstrating how the community could extend the ontology over time. The ontology can be accessed at **https://ontology.darkpatternsresearchandimpact.com**, which includes a current state of the ontology and community-vetted changes over time that follow this process in a public, deliberative manner. Initial and future iterations of the ontology will be versioned and include a version history for citation accuracy.

Extending the Ontology			
Name	**Definition from the Sources**	**Mapping to Ontology**	**Level**
Linguistic Dead-End [30]	"[D]esign patterns wherein language use prevents or makes it very difficult for the user to understand crucial functionality [...]".	*Language Inaccessibility*	extends meso-level
Untranslation [30]	"[D]esign patterns in which part or all of the app is in a language unfamiliar to the people using it, even if the app is stated as available in the local language in the store".	*Wrong Language*	extends low-level
Alphabet Soup [30]	"[D]esign pattern language use prevents or makes it very difficult for the user to understand crucial functionality [...]".	*Language Inaccessibility*	new low-level
Extraneous Badges [29]	"[D]esign elements — often tiny, brightly colored circles—that visually highlight UI elements that require immediate user attention".	*Aesthetic Manipulation*	new low-level
Account Deletion Roadblocks [29]	"*Unclear deactivation/deletion options* covers cases where a service insufficiently communicates what will happen if a person deactivates or deletes their account."	*Roach Motel*	new low-level
	"*Time-Delayed Account Deletion* covers cases where a service will only initiate the account deletion process after a cool-off period, rather than instantaneously."	*Roach Motel*	new low-level
Engaging Strategies [43]	"[D]ark patterns where the goal is to keep users occupied and entertained for as long as possible".	*Social Engineering*	extends high-level
Governing Strategies [43]	Dark patterns "that navigate users' decision-making towards the designers' and/or platform providers' goals".	*Obstruction*	extends high-level
Labyrinthine Navigation [43]	"[N]ested interfaces that are easy to get lost in, disabling users from choosing preferred settings".	*Privacy Maze*	extends low-level

Table 1: This table presents an overview of selected dark patterns from Hidaka et al. [30], Gunawan et al. [29], and Mildner et al. [43] to demonstrate extending the dark pattern ontology.

6.1 Dark Patterns In Japanese Apps

Hidaka et al. [30] studied dark patterns in Japanese apps and identified two dark pattern types—*Unstranslation* and *Alphabet Soup*—which are sub-types of a novel *Linguistic Dead-End* dark pattern. They specifically motivated their work as one of the first studies of dark patterns in a non-Western context. We closely evaluated the authors' definition of *Linguistic Dead-End*, where the use of a foreign language hinders users from understanding the consequences of their interactions. When comparing these three patterns to our ontology, the high-level pattern *Linguistic Dead-End* appears to fit within the existing meso-level dark pattern *Language Inaccessibility* while extending its coverage. The remaining two low-level patterns, *Untranslation* and *Alphabet Soup*, can then be nested as two low-level types underneath the same meso-level dark pattern, with *Untranslation* mapping to and extending the existing *Wrong Language* dark pattern and *Alphabet Soup* forming a new low-level pattern. In this case, the three dark patterns extend and further support a distinct area of the ontology, demonstrating how novel contexts help to usefully supplement existing dark patterns and identify new low-level means of execution. Additionally, this study demonstrates that dark patterns exist across multiple cultures and areas of the world, but may take different forms depending on local design norms.

6.2 Contextual Dark Patterns in Different Screen Modalities

Gunawan et al. [29] investigated the presence of dark patterns across different screen modalities, describing eight novel dark pattern types which limit the choices of users depending on the device used. In the provided definitions for each proposed dark pattern, the authors included links to previously defined dark patterns—linking these patterns to elements of the ontology, thus providing an easy

mapping path. The *Extraneous Badges* dark pattern, for example, is indicated as related to *Aesthetic Manipulation* [24] as a form of *Interface Interference*, and would result in this dark pattern being included as a new low-level type in the ontology. Similarly, using the authors' definitions and identification of mapping in the paper text, *Account Deletion Roadblocks* could extend *Roach Motel* through two specific new low-level types focusing variously on insufficient communication and time delay: *Unclear Deactivation/Deletion Options* and *Time-Delayed Account Deletion*. These examples illustrate how contextual and situational links to previously defined dark patterns support the ontology, describing specific situations that strengthen established dark patterns and identify new low-level means of execution.

6.3 Domain-Specific Dark Patterns in Social Media Applications

Mildner et al. [43] investigated dark patterns on social media platforms, proposing five dark patterns across two strategies. As with Hidaka et al., the granularity of their definitions implies a mapping on multiple levels of the ontology. We began by drawing from the authors' definitions of *Engaging Strategies* and *Governing Strategies*. The authors describe the aim of *Engaging Strategies* as entertaining users for as long as possible, related to *Attention Capture* [44], which is already included in the ontology as a meso-level pattern under *Forced Action*. However, some elements of the original definition (e.g., occupying and entertaining) fit more closely within concepts of *Social Engineering*. Similarly, *Governing Strategies* can be partially linked to multiple patterns in the ontology. For example, as the authors originally suggest, the strategy can be enabled through *Interface Interference*. However, *Governing Strategies* also offers a high-level focus to inspect *Obstruction* with *Labyrinthine Navigation*, presenting an interesting adaption of *Privacy Maze* already present in the ontology. These examples indicate how the authors make their dark pattern types distinct from prior ones, functioning as a lens that might invite reinspection of dark patterns in the ontology and perhaps indicate opportunities for further development of low-level patterns.

6.4 Dark Patterns in Legislation and Case Law

An alignment between legislation, the ontology, and case law shows that it could also be a robust and reliable artifact for regulators and policy makers to use in their compliance monitoring and enforcement actions.

Mapping the ontology to case law Dark patterns have been detected in regulatory cases by enforcers, such as Data Protection Authorities (DPAs) and Consumer Protection Authorities, for more than a decade [15, 21]. However few cases explicitly designate dark patterns as such.[7] Decisions analyse several practices that are related to dark patterns, but without qualifying each practice into a concrete granular type of dark pattern. Current case law descriptions of the use of dark patterns often report infringements only at a general level, but without qualifying each practice as a concrete type of dark pattern [50]. In doing so, case law could miss lower-level granularity that may translate across domains. A recent example shows that a EU regulator, the Italian DPA, used the concept of dark patterns related to certain consent practices for the first time in an official EU legal decision [34]. By mapping case law to the ontology, regulators can gain additional knowledge identifying where dark patterns practices at multiple levels and in multiple combinations are at play, and were deemed to be illegal per jurisdiction [36], enhancing legal certainty about dark patterns practices. For example, the EU Court of Justice has ruled that the practice called "pre-selection" violates the GDPR [9], which maps to the meso-level dark pattern *"Bad Defaults"* in our ontology.

Further, the ontology has the potential to support enforcement decisions since it can test and confirm the traceability of concrete dark patterns-related practices. For instance, the Italian Data Protection Authority has already added the keyword "dark pattern" to the available tags of their online database[8]—a useful effort that should be extended to official and unofficial searchable databases of enforcers' decisions. Connecting case law to multiple levels of dark patterns in our proposed ontology has the potential to inform enforcers of different jurisdictions in the EU/US and reduce the risks of gaps or overlaps.

Mapping the ontology to legislation The proposed ontology can also help regulators across different jurisdictions to understand relationships between different definitions of dark patterns, including high-, meso- and low level dark patterns, including when such definitions map to existing and upcoming legislation. The recent EU Digital Service Act (DSA)[14, Art.25(3)(b), recital 67] explicitly prohibits user manipulation and specifies that further guidelines will be given on a specific practice, where "repeatedly requesting that the recipient of the service make a choice where that choice has already been made, especially by presenting pop-ups that interfere with the user experience"; this example maps well to the proposed *Nagging* dark pattern in our ontology. Because new legislation, such as the DSA[14], Data Market Act (DMA)[13], Data Act [12], and California CPRA [11] contain dark patterns specific prohibitions, we believe the proposed ontology has the capability to ensure a precise mapping between the concepts of dark patterns in research literature and the legally-binding provisions. When the concepts of the ontology are mapped to a legal concept, then it is easier for regulators to link a specific dark pattern to a concrete binding legislative provision. Consequently, the ontology will help to conclude the normative value of such practice—whether a specific dark pattern is illegal or legal—and what relevant obligations and rights are derived from the law and must be enforced. If regulators and policy-makers across jurisdictions rely on the same definitions of dark patterns, this can assure an easier re-use of case law for future legal cases.

7 USING THE ONTOLOGY TO SUPPORT TRANSDISCIPLINARY ENGAGEMENT

In this ontology, we seek to synthesize and harmonize existing academic and regulatory taxonomies while adding useful and consistent structure to allow for other stakeholders to build upon and derive benefit from a shared description of dark patterns knowledge. This paper lays the foundation for shared action, which includes many different stakeholders with differing aims. In this section,

[7] Case law and legal frameworks have recently been added to the https://deceptive.design site, which includes mappings to specific dark patterns [2].

[8] https://www.garanteprivacy.it/temi/internet-e-nuove-tecnologie/dark-pattern

we outline key opportunities for future transdisciplinary engagement, identifying opportunities for scholars to continue building knowledge about dark patterns and their harms, for regulators and other enforcement agencies to better detect and thus sanction dark patterns, and for legal scholars and legislators to address current and future consequences of dark patterns that can inform further action.

7.1 Challenges in Evolving the Ontology

Not all of our mappings were clear-cut and some may be productively extended or disputed in future versions of this ontology. Through dialogue, we sought to locate existing patterns within our ontology based on our best understanding of the pattern as described by its name and definition in the source taxonomy. One challenge we faced was that some combinations of patterns have evolved over time. For instance, Mathur et al.'s [40] high-level pattern "social proof" originated with two sub-patterns, "activity messages" and "testimonials." Later, the FTC created new low-level patterns, introducing "endorsements" (we bundled it with testimonials as one low-level pattern) and more specific types of endorsement or testimonials (e.g., "deceptive celebrity endorsements," "false activity messages"). Future work could identify the most useful level of abstraction for these patterns.

Additionally, the use of novel names for patterns (particularly by the EDPB and CMA) or the use of patterns in specific contexts (e.g., e-commerce, social media) caused us to consider both the presence of granular low-level patterns and the relation of these low-level patterns to inferred meso-level patterns. In particular, the use of novel names for patterns types and definitions was a challenge from an analytic perspective, resulting in: i) instances where a wholly new pattern was introduced (e.g., CMA's "information overload" which we leveraged to create a new meso-level pattern of "choice overload"); ii) instances where a new high-level strategy was highly similar to an existing high-level strategy (e.g., EDPB's "skipping" which we subsumed within "sneaking"); and iii) instances where existing patterns included both a generalizable pattern and domain-specific information which may need to be captured in specific low-level patterns in future work (e.g., EDPB's "left in the dark" is a form of "hidden information" but implies specific low-level patterns that are specific to data protection).

These observed challenges point towards the value of a shared ontology that includes a consistent vocabulary, but also points to opportunities to generate more specific knowledge that is linked to particular contexts and technologies. For instance, low-level patterns could be tagged based on how well they relate to specific contexts (e.g., e-commerce, social media), technologies (e.g., CUIs, VR/AR, robots), or application domains (e.g., health, travel) as indicated by a recent systematic review of dark patterns literature [22].

Finally, formal evaluation of the definitions and ontology structure we have proposed will strengthen our understanding of how various stakeholders consider, interpret, and use the ontology to support their work—within and across technology contexts. For instance, the language specificity demanded by a legal or regulatory professional from a given definition within the ontology may require different kinds of analytic precision as compared to the generative or evaluative use of the same definitions by a designer performing an audit of dark patterns on digital systems for their company. Future work should address both the utility and the rigor of various components of our ontology for differing purposes, including expert evaluation, gathering of evidence for legal and regulatory action, operationalization of dark patterns for social science research, and use by designers to avoid inscribing dark patterns into their design work.

7.2 Activating Transdisciplinary Pathways

As we have outlined, work relating to dark patterns has connected many different disciplinary communities toward shared goals, including social scientists studying the presence and harms of dark patterns, legal scholars linking instances of dark patterns to relevant consumer protection or data protection legal frameworks, legislators targeting specific legal provisions about dark patterns to support new obligations and/or future sanctions, and regulators detecting legal violations related to dark patterns to support enforcement sanctions. We consider multiple opportunities for collaboration within and across these stakeholder groups:

- **Social Scientists** Scientists studying dark patterns can use the ontology to better map the impact triggered by certain dark patterns in concrete contexts in ways that support shared knowledge building and reduce duplication. This approach has been applied for specific low-level patterns by various empirical studies that evaluated the impact of dark pattern design on the outcome of users' consent decisions [45], but could be scaled up substantially using the ontology as a means of producing and sharing these mappings.
- **Social Scientists + Computer Scientists** The detection of dark patterns could also be more robustly supported by our ontology, with our assertion that low-level patterns show the most promise in being detectable. Existing detection efforts (e.g., [5, 8, 35, 40, 46, 52–54, 57]) have shown that higher-level patterns are difficult or impossible to detect at scale due to their abstract nature that requires interpretation, while low-level UI elements with discrete and known qualities (e.g., cookie consent banners, elements of the checkout process) are more detectable using software tools for automated detection. Our ontology of low-level patterns and gaps creates a foundation for future detection efforts, allowing computer science scholars to focus on pattern types which are most likely to be detectable and measurable.
- **Social Scientists + Regulators** Bielova et al. [3] have recently compared the results of such empirical studies and designs recommended by EU regulators and found multiple gaps and contradictions relating to instances of dark patterns, showing that empirical studies bring important insights not only in the research community but also for the regulators and policy-makers. This effort demonstrates an opportunity for regulators and social scientists to work more closely—commissioning studies where user experience of dark patterns is unknown or unclear (particularly with relation to causal mechanisms) while deprioritizing studies

that address design choices that are already illegal under statute.

- **Social Scientists + Legal Scholars** The ontology can be extended to consider potential harms in relation to specific dark patterns types [28]. For example, the meso-level dark pattern *Nagging* can arguably trigger "attentional theft," thus harming consumer welfare, and can lead to indirect harms such as increased vulnerability to privacy violations, and finally, to anti-competitive harms [32]. A mapping of harms to specific types of dark patterns in the ontology may support connections to avenues for legal remedies, as well as aid in identifying areas where additional research is needed.
- **Legal Scholars + Regulators** The ontology may also be extended to refer to concrete enforcement cases already consolidated in a database of dark patterns case law, such as those on Brignull's updated site [2]. This will allow for case law to inform future legal sanctions, identify which elements of the ontology connect to existing legal frameworks, and lay the groundwork for future legislative action to allow for sanctioning of novel patterns that are not well addressed through existing laws.

8 CONCLUSION

To support the development of a shared language of dark patterns, in this paper we present our analysis of ten existing regulatory and academic taxonomies of dark patterns and propose a three-level ontology with standardized definitions for 65 synthesized dark pattern types across low-, meso-, and high-level patterns. Building on our analysis, future scholars, regulators, and legal professionals can benefit from our hierarchical organization of dark patterns types to indicate links to existing and similar concepts. This description encourages the establishment of provenance in future work, allowing scholars and regulators to identify pattern types and their origins and provide an audit trail to connect specific contextually-bound instances with broader categorizations. This ontology creates a foundation for a shared and reusable knowledge source, allowing many stakeholders to work together in building a shared, explicit and precise conceptualization of what is already known in the literature and which can be further refined and extended. Finally, we illustrate how this ontology can support translational research and regulatory action, by extending the ontology from three contemporary studies defining dark patterns from domain and context-specific areas, as well as ontology extension to map legislation and case law.

ACKNOWLEDGMENTS

This work is funded in part by the National Science Foundation under Grant No. 1909714 and the ANR 22-PECY-0002 IPOP (Interdisciplinary Project on Privacy) project of the Cybersecurity PEPR. The research of this work was partially supported by the Klaus Tschira Stiftung gGmbH.

REFERENCES

[1] 2023. Central Consumer Protection Authority issues 'Guidelines for Prevention and Regulation of Dark Patterns, 2023' for prevention and regulation of dark patterns listing 13 specified dark patterns. https://pib.gov.in/PressReleaseIframePage.aspx?PRID=1983994

[2] 2023. Deceptive patterns - Legal Cases. https://www.deceptive.design/cases Accessed: 2023-9-14.

[3] Nataliia Bielova, Cristiana Santos, and Colin M Gray. 2024. Two worlds apart! Closing the gap between regulating EU consent and user studies. *Harvard Journal of Law & Technology* 37 (2024).

[4] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns. *Proceedings on Privacy Enhancing Technologies* 2016, 4 (2016). https://doi.org/10.1515/popets-2016-0038

[5] Ahmed Bouhoula, Karel Kubicek, Amit Zac, Carlos Cotrini, and David Basin. 2023. Automated, Large-Scale Analysis of Cookie Notice Compliance. In *USENIX Security Symposium*.

[6] Harry Brignull. 2018. Deceptive Patterns: User Interfaces Designed to Trick People. http://darkpatterns.org/

[7] Harry Brignull. 2023. Deceptive Patterns. https://www.deceptive.design

[8] Jieshan Chen, Jiamou Sun, Sidong Feng, Zhenchang Xing, Qinghua Lu, Xiwei Xu, and Chunyang Chen. 2023. Unveiling the Tricks: Automated Detection of Dark Patterns in Mobile Applications *(UIST '23 Adjunct)*. Association for Computing Machinery, San Francisco, CA, USA, 1–20. http://arxiv.org/abs/2308.05898 arXiv:2308.05898 [cs].

[9] CJEU-Planet49-19 2019. Judgment in Case C-673/17 Bundesverband der Verbraucherzentralen und Verbraucherverbände – Verbraucherzentrale Bundesverband eV v Planet49 GmbH. http://curia.europa.eu/juris/documents.jsf?num=C-673/17.

[10] CMA2022 2022. *Evidence review of Online Choice Architecture and consumer and competition harm.* Technical Report. https://www.gov.uk/government/publications/online-choice-architecture-how-digital-design-can-harm-competition-and-consumers/evidence-review-of-online-choice-architecture-and-consumer-and-competition-harm Accessed: 2022-4-13.

[11] CPRA 2022. California Privacy Rights Act. https://cppa.ca.gov/meetings/materials/20220608_item3.pdf

[12] Data-Act-proposal 2022. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on harmonised rules on fair access to and use of data (Data Act). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2022%3A68%3AFIN

[13] DMA 2022. Digital Markets Act - Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act) (Text with EEA relevance). http://data.europa.eu/eli/reg/2022/1925/oj

[14] DSA2022 2022. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act).

[15] 2022. *Behavioural study on unfair commercial practices in the digital environment : dark patterns and manipulative personalisation : final report.* Publications Office of the European Union. https://op.europa.eu/en/publication-detail/-/publication/606365bc-d58b-11ec-a95f-01aa75ed71a1/language-en/format-PDF/source-257599418

[16] European Data Protection Board. 2022. *Guidelines 3/2022 on Dark patterns in social media platform interfaces: How to recognise and avoid them.* Technical Report Version 1.0. https://edpb.europa.eu/system/files/2022-03/edpb_03-2022_guidelines_on_dark_patterns_in_social_media_platform_interfaces_en.pdf

[17] European Data Protection Board. 2023. *Guidelines 3/2022 on Dark patterns in social media platform interfaces: How to recognise and avoid them.* Technical Report Version 2.0. https://edpb.europa.eu/system/files/2023-02/edpb_03-2022_guidelines_on_deceptive_design_patterns_in_social_media_platform_interfaces_v2_en_0.pdf

[18] Frederico Fonseca. 2007. The double role of ontologies in information science research. *Journal of the American Society for Information Science and Technology* 58, 6 (April 2007), 786–793. https://doi.org/10.1002/asi.20565

[19] French DPA (CNIL). 2021. Deliberation of the restricted committee No. SAN-2021-023 of 31 December 2021 concerning GOOGLE LLC and GOOGLE IRELAND LIMITED. https://www.cnil.fr/sites/default/files/atoms/files/deliberation_of_the_restricted_committee_no._san-2021-023_of_31_december_2021_concerning_google_llc_and_google_ireland_limited.pdf.

[20] French DPA (CNIL). 2022. Deliberation of the restricted committee No. SAN-2021-024 of 31 December 2021 concerning FACEBOOK IRELAND LIMITED. https://www.cnil.fr/sites/default/files/atoms/files/deliberation_of_the_restricted_committee_no._san-2021-024_of_31_december_2021_concerning_facebook_ireland_limited.pdf.

[21] FTC2022 2022. *Bringing Dark Patterns to Light Staff Report.* Technical Report. Federal Trade Commission. https://www.ftc.gov/system/files/ftc_gov/pdf/P214800%20Dark%20Patterns%20Report%209.14.2022%20-%20FINAL.pdf

[22] Colin M Gray, Lorena Sánchez Chamorro, Ike Obi, and Ja-Nae Duane. 2023. Mapping the Landscape of Dark Patterns Scholarship: A Systematic Literature Review. In *Designing Interactive Systems Conference (DIS Companion '23)* (Pittsburgh, PA, USA), Vol. 1. Association for Computing Machinery. https://doi.org/10.1145/3563703.3596635

[23] Colin M Gray, Johanna Gunawan, René Schäfer, Nataliia Bielova, Lorena Sánchez Chamorro, Katie Seaborn, Thomas Mildner, and Hauke Sandhaus. 2024. Mobilizing Research and Regulatory Action on Dark Patterns and Deceptive Design Practices. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA'24)*. Association for Computing Machinery. https://doi.org/10.1145/3613905.3636310

[24] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. 2018. The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. dl.acm.org, New York, NY, USA, 534:1–534:14. https://doi.org/10.1145/3173574.3174108

[25] Colin M Gray, Cristiana Santos, and Nataliia Bielova. 2023. Towards a Preliminary Ontology of Dark Patterns Knowledge. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. https://doi.org/10.1145/3544549.3585676

[26] Colin M Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damian Clifford. 2021. Dark Patterns and the Legal Requirements of Consent Banners: An Interaction Criticism Perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI'21)*. ACM Press. https://doi.org/10.1145/3411764.3445779

[27] Colin M Gray, Cristiana Santos, Nicole Tong, Thomas Mildner, Arianna Rossi, Johanna Gunawan, and Caroline Sinders. 2023. Dark Patterns and the Emerging Threats of Deceptive Design Practices. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. https://doi.org/10.1145/3544549.3583173

[28] Gunawan, Santos, and Kamara. 2022. Redress for Dark Patterns Privacy Harms? A Case Study on Consent Interactions. *2022 ACM Symposium on* (2022). https://johannagunawan.com/assets/pdf/gunawan-22-cslaw.pdf

[29] Johanna Gunawan, Amogh Pradeep, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2021. A Comparative Study of Dark Patterns Across Web and Mobile Modalities. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (Oct. 2021), 1–29. https://doi.org/10.1145/3479521

[30] Shun Hidaka, Sota Kobuki, Mizuki Watanabe, and Katie Seaborn. 2023. Linguistic Dead-Ends and Alphabet Soup: Finding Dark Patterns in Japanese Apps. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–13. https://doi.org/10.1145/3544548.3580942

[31] Hsiu-Fang Hsieh and Sarah E Shannon. 2005. Three approaches to qualitative content analysis. *Qualitative health research* 15, 9 (Nov. 2005), 1277–1288. https://doi.org/10.1177/1049732305276687

[32] Alison Hung. 2021. KEEPING CONSUMERS IN THE DARK: ADDRESSING "NAGGING" CONCERNS AND INJURY. *Columbia Law Review* 121, 8 (2021), 2483–2520. https://www.jstor.org/stable/27093855

[33] india-2023-ek 2023. Draft Guidelines for Prevention and Regulation of Dark Patterns. https://consumeraffairs.nic.in/sites/default/files/file-uploads/latestnews/Draft%20Guidelines%20for%20Prevention%20and%20Regulation%20of%20Dark%20Patterns%202023.pdf

[34] ItalianDPAvsEdiscom2023 2023. Provvedimento prescrittivo e sanzionatorio nei confronti di Ediscom S.p.A. - 23 febbraio 2023 [9870014]. https://www.garanteprivacy.it/web/guest/home/docweb/-/docweb-display/docweb/9870014

[35] Simon Koch, Benjamin Altpeter, and Martin Johns. 2023. The OK is not enough: Large Scale Study of Consent Dialogs in Smartphone Applications. In *USENIX Security Symposium*.

[36] Mark Leiser. 2020. *'Dark Patterns': the case for regulatory pluralism.* LawArXiv ea5n2. Center for Open Science. https://doi.org/10.31219/osf.io/ea5n2

[37] Jamie Luguri and Lior Jacob Strahilevitz. 2021. Shining a Light on Dark Patterns. *Journal of Legal Analysis* 13, 1 (March 2021), 43–109. https://doi.org/10.1093/jla/laaa006

[38] Kai Lukoff, Alexis Hiniker, Colin M Gray, Arunesh Mathur, and Shruthi Sai Chivukula. 2021. What can CHI do about dark patterns?. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama Japan). ACM, New York, NY, USA. https://doi.org/10.1145/3411763.3441360

[39] Luxembourg DPA. 2021. Decision regarding Amazon Europe Core S.À RL. https://cnpd.public.lu/fr/actualites/international/2021/08/decision-amazon-2.html.

[40] Arunesh Mathur, Gunes Acar, Michael J Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), Article No. 81. https://doi.org/10.1145/3359183

[41] Arunesh Mathur, Mihir Kshirsagar, and Jonathan Mayer. 2021. What Makes a Dark Pattern... Dark? Design Attributes, Normative Considerations, and Measurement Methods. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3411764.3445610

[42] Thomas Mildner and Gian-Luca Savino. 2021. Ethical User Interfaces: Exploring the Effects of Dark Patterns on Facebook. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21, Article 464)*. Association for Computing Machinery, New York, NY, USA, 1–7. https://doi.org/10.1145/3411763.3451659

[43] Thomas Mildner, Gian-Luca Savino, Philip R. Doyle, Benjamin R. Cowan, and Rainer Malaka. 2023. About Engaging and Governing Strategies: A Thematic Analysis of Dark Patterns in Social Networking Services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–15. https://doi.org/10.1145/3544548.3580695

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

[45] Nataliia Bielova. 2023. A survey of user studies as evidence for dark patterns in consent banners. https://backoffice.cnil.fr/sites/default/files/atoms/files/full_2022-12-02_v2.pdf, accessed on 7 September 2023..

[46] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376321

[47] Ikechukwu Obi, Colin M Gray, Shruthi Sai Chivukula, Ja-Nae Duane, Janna Johns, Matthew Will, Ziqing Li, and Thomas Carlock. 2022. Let's Talk About Socio-Technical Angst: Tracing the History and Evolution of Dark Patterns on Twitter from 2010-2021. (July 2022). arXiv:2207.10563 [cs.SI] http://arxiv.org/abs/2207.10563

[48] OECD. 2022. *Dark commercial patterns*. Technical Report. https://doi.org/10.1787/44f5e846-en

[49] Press-Release2022-fq 2022. Press Release: AG Racine announces Google must pay $9.5 million for using "dark patterns" and deceptive location tracking practices that invade users' privacy. https://thedcline.org/2022/12/30/press-release-ag-racine-announces-google-must-pay-9-5-million-for-using-dark-patterns-and-deceptive-location-tracking-practices-that-invade-users-privacy/. https://thedcline.org/2022/12/30/press-release-ag-racine-announces-google-must-pay-9-5-million-for-using-dark-patterns-and-deceptive-location-tracking-practices-that-invade-users-privacy/ Accessed: 2022-12-31.

[50] Cristiana Santos and Arianna Rossi. 2023. The emergence of dark patterns as a legal concept in case law. https://policyreview.info/articles/news/emergence-of-dark-patterns-as-a-legal-concept

[51] Brennan Schaffner, Neha A Lingareddy, and Marshini Chetty. 2022. Understanding Account Deletion and Relevant Dark Patterns on Social Media. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (Nov. 2022), 1–43. https://doi.org/10.1145/3555142

[52] Than Htut Soe, Cristiana Teixeira Santos, and Marija Slavkovik. 2022. Automated detection of dark patterns in cookie banners: how to do it poorly and why it is hard to do it any other way. (April 2022). arXiv:2204.11836 [cs.LG] http://arxiv.org/abs/2204.11836

[53] Ioannis Stavrakakis, Andrea Curley, Dympna O'Sullivan, Damian Gordon, and Brendan Tierney. 2021. A Framework of Web-Based Dark Patterns that can be Detected Manually or Automatically. (2021). https://doi.org/10.21427/20g8-d176

[54] Marieke Van Hofslot, Almila Akdag Salah, Albert Gatt, and Cristiana Santos. 2022. Automatic Classification of Legal Violations in Cookie Banner Texts. In *Proceedings of the Natural Legal Language Processing Workshop 2022*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Hybrid), 287–295. https://aclanthology.org/2022.nllp-1.27

[55] Jess Weatherbed. 2022. Google is paying a $85m settlement to Arizona to end user-tracking suit. https://www.theverge.com/2022/10/5/23389331/google-settlement-arizona-user-tracking-privacy-suit. https://www.theverge.com/2022/10/5/23389331/google-settlement-arizona-user-tracking-privacy-suit Accessed: 2023-1-4.

[56] Shoshana Wodinsky. 2022. The 'dark patterns' in Fortnite that led to the largest FTC penalties ever. https://www.marketwatch.com/story/the-dark-patterns-in-fortnite-that-led-to-the-largest-ftc-penalties-ever-11671488228. https://www.marketwatch.com/story/the-dark-patterns-in-fortnite-that-led-to-the-largest-ftc-penalties-ever-11671488228 Accessed: 2022-12-20.

[57] Yuki Yada, Jiaying Feng, Tsuneo Matsumoto, Nao Fukushima, Fuyuko Kido, and Hayato Yamana. 2022. Dark patterns in e-commerce: a dataset and its baseline evaluations. In *2022 IEEE International Conference on Big Data (Big Data)*. 3015–3022. https://doi.org/10.1109/BigData55660.2022.10020800

[58] José P Zagal, Staffan Björk, and Chris Lewis. 2013. Dark Patterns in the Design of Games. In *Foundations of Digital Games*. http://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1043332&dswid=1018

A FINAL ONTOLOGY DEFINITIONS

- **Sneaking** is a strategy which hides, disguises, or delays the disclosure of important information that, if made available to users, would cause a user to unintentionally take an action they would likely object to.
  - **Bait and Switch** subverts the user's expectation that their choice will result in a desired action, instead leading to an unexpected, undesirable outcome.
    - **Disguised Ads** *Bait and Switch* and use *Sneaking* to style interface elements so they are not clearly marked as an advertisement or other biased source. As a result, users are induced into clicking on the interface element because they assume that it is a relevant and salient interaction, leading to unwitting interaction with advertising content.
  - **Hiding Information** subverts the user's expectation that all relevant information to make an informed choice will be available to them, instead hiding information or delaying the disclosure of information until later in the user journey that may have led to them making another choice.
    - **Sneak into Basket** *Hides Information* and uses *Sneaking* to add unwanted items to a user's shopping cart without their consent. As a result, a user assumes that only the items they explicitly added to their cart will be purchased, leading to unintentional purchase of additional items.
    - **Drip Pricing, Hidden Costs, or Partitioned Pricing** *Hides Information* and uses *Sneaking* to reveal new charges or costs, present only partial price components, or otherwise delay revealing the full price of a product or service through late or incomplete disclosure. As a result, the user is misled about the total or complete price of the product or service, leading to them to make a purchase decision after they have expended effort on false pretenses.
    - **Reference Pricing** *Hides Information* and uses *Sneaking* to include a misleading or inaccurate price for a product or service that makes a discounted price appear more attractive. As a result, the user is misled into believing that the price they pay is discounted, leading them to make a decision to purchase a product or service on false pretenses.
  - **(De)contextualizing Cues** subverts the user's expectation that provided information will guide the user to making an informed choice, instead confusing the user and/or preventing them from locating relevant information due to the context where information is presented.
    - **Conflicting Information** uses *(De)contextualizing Cues* and *Sneaking* to include two or more sources of information that conflict with each other. As a result, the user is unsure what the consequences of their actions will be and will be more likely to accept default settings that may not be in their best interest.
    - **Information without context** uses *(De)contextualizing Cues* and *Sneaking* to alter the relevant information or user controls to limit discoverability. As a result, the user is unlikely to find the information or action possibility they are interested in.
- **Obstruction** is a strategy which impedes a user's task flow, making an interaction more difficult than it inherently needs to be, dissuading a user from taking an action.
  - **Roach Motel** subverts the user's expectation that an action will be as easy to reverse as it is to make, instead creating a situation that is easy to get into, but difficult to get out of.
    - **Immortal Accounts** create a *Roach Motel* and use *Obstruction* to make it difficult or impossible to delete a user account once it has been created. As a result, the user may create an account or share data with the false assumption that they can later delete this information, even though that account and/or data are then unable to be removed by the user.
    - **Dead Ends** create a *Roach Motel* and use *Obstruction* to prevent users from finding information through inactive links or redirections that limit or completely prevent the display of relevant information. As a result, the user may seek to find relevant information or action possibilities but instead be left unable to achieve their goal.
  - **Creating Barriers** subverts the user's expectation that relevant user tasks will be supported by the interface, instead preventing, abstracting, or otherwise complicating a user task to disincentive user action.
    - **Price Comparison Prevention** *Creates Barriers* and uses *Obstruction* by excluding relevant information, limiting the ability of a user to copy/paste, or otherwise inhibiting a user from comparing prices across two or more vendors. As a result, the user cannot make an informed decision about where to buy a product or service.
    - **Intermediate Currencies** *Create Barriers* and use *Obstruction* to hide the true cost of a product or service by requiring the user to spend real money to purchase a virtual currency that is then used to purchase a product or service. As a result, the user is unable to easily ascertain the true monetary cost of a product or service, leading them to make an uninformed purchase decision based on an obscured cost.
  - **Adding Steps** subverts the user's expectation that a task will take as few steps as technologically needed, instead creating additional points of unnecessary but required user interaction to perform a task.
    - **Privacy Mazes** *Add Steps* and use *Obstruction* to require a user to navigate through many pages to obtain relevant information or control without a comprehensive and exhaustive overview. As a result, the user is prevented from easily discovering relevant information or action possibilities, leaving them unable to make informed decisions regarding their privacy.
- **Interface Interference** is a strategy which privileges specific actions over others through manipulation of the user interface, thereby confusing the user or limiting discoverability of relevant action possibilities.

- **Manipulating Choice Architecture** subverts the user's expectation that the options presented will support their desired goal, instead including an order or structure of options that makes another outcome more likely.
  - **False Hierarchy** *Manipulates the Choice Architecture,* using *Interface Interference* to give one or more options visual or interactive prominence over others, particularly where items should be in parallel rather than hierarchical. As a result, the user may misunderstand or be unable to accurately compare their options, making a selection based on a false or incomplete choice architecture.
  - **Visual Prominence** *Manipulates the Choice Architecture,* using *Interface Interference* to place an element relevant to user goals in visual competition with a more distracting and prominent element. As a result, the user may forget about or be distracted from their original goal, even if that goal was their primary intent.
  - **Bundling** *Manipulates the Choice Architecture,* using *Interface Interference* to group two or more products or services in a single package at a special price. As a result, the user may incorrectly assume that these items must be purchased as a bundle or be unaware of the unbundled price for the component elements, possibly leading to an uninformed purchasing decision.
  - **Pressured Selling** *Manipulates the Choice Architecture,* using *Interface Interference* to preselect or use visual prominence to focus user attention on more expensive product options . As a result, the user may be unaware that a lower price is available or even desirable for their needs , steering the user into making a more expensive product selection than they otherwise would have.
- **Bad Defaults** subverts the user's expectation that default settings will be in their best interest, instead requiring users to take active steps to change settings that may cause harm or unintentional disclosure of information.
- **Emotional or Sensory Manipulation** subverts the user's expectation that the design of the site will allow them to achieve their goal without manipulation, instead altering the language, style, color, or other design elements to evoke an emotion or manipulate the senses in order to persuade the user into a particular action.
  - **Cuteness** uses *Emotional or Sensory Manipulation* and *Interface Interference* to embed attractive cues in the design of a robot interface or form factor. As a result, a user may place undue trust in the robot, leading the user to inaccurately or incompletely assess the risks of interacting with the robot.
  - **Positive or Negative Framing** uses *Emotional or Sensory Manipulation* and *Interface Interference* to visually obscure, distract, or persuade a user from important information they need to achieve their goal. As a result, the user may assume that the system is providing equal access to relevant information, leading the user to be distracted by positive or negative aesthetic cues that distract them from important information or action possibilities or otherwise convince them to pursue a different goal.
- **Trick Questions** subvert the user's expectation that prompts will be written in a straightforward and intelligible manner, instead using confusing wording, double negatives, or otherwise leading language or interface cues to manipulate a user's choice.
- **Choice Overload** subverts the user's expectation that the choices they make should be understandable and comparable, instead providing too many options to compare or encouraging users to overlook relevant information due to the volume of choices provided.
- **Hidden Information** subverts the user's expectation that relevant information will be made accessible and visible, instead disguising relevant information or framing it as irrelevant.
- **Language Inaccessibility** subverts the user's expectation that guidance will be provided in a way that is understandable and intelligible, instead using unnecessarily complex language or a language not spoken by the user to decrease the likelihood the user will make an informed choice.
  - **Wrong Language** leverages *Language Accessibility*, using *Interface Interference* to provide important information in a different language than the official language of the country where users live. As a result, the user will not have access to relevant information about their interaction with the system and their ability to choose, leading to uninformed decisions.
  - **Complex Language** leverages *Language Accessibility*, using *Interface Interference* to make information difficult to understand by using obscure word choices and/or sentence structure. As a result, the user will not be able to comprehend relevant information about their interaction with the system and their ability to choose, leading to uninformed decisions.
- **Feedforward Ambiguity** subverts the user's expectation that their choice will be likely to result in an action they can predict, instead providing a discrepancy between information and actions available to users that results in an outcome that is different from what the user expects.

- **Forced Action** is a strategy which requires users to knowingly or unknowingly perform an additional and/or tangential action or information to access (or continue to access) specific functionality, preventing them from continuing their interaction with a system without performing that action.
  - **Nagging** subverts the user's expectation that they have rational control over the interaction they make with a system, instead distracting the user from a desired task the user is focusing on to induce an action or make a decision the user does not want to make by repeatedly interrupting the user during normal interaction.
  - **Forced Continuity** subverts the user's expectation that a subscription created in the past will not auto-renew or otherwise continue in the future, instead causing undesired charges, difficulty to cancel, or lack of awareness that a subscription is still active.
  - **Forced Registration** subverts the user's expectation that they can complete an action without registering or creating an account, instead tricking them into thinking that registration is required, often resulting in the sharing of unneeded personal data.

  - **Forced Communication or Disclosure** subverts the user's expectation that a system will only request information needed to complete their desired goals, instead tricking them into sharing more information about themselves or using their information for purposes that they do not desire.
    - **Privacy Zuckering** uses *Forced Communication or Disclosure* as a type of *Forced Action* to trick users into sharing more information about themselves than they intend to or would agree to if fully informed. As a result, the user assumes that information they are requested to provide is vital for use of the service, even while this information is used or sold for other purposes.
    - **Friend Spam** uses *Forced Communication or Disclosure* as a type of *Forced Action* to collect information about other users through extractive means that results in unwanted contact from the service. As a result, the user assumes that information about their friends or social network is vital for use of the service, even while this information is used to spam other users.
    - **Address Book Leeching** uses *Forced Communication or Disclosure* as a type of *Forced Action* to collect information about other users through extractive means, which are often hidden to the user and/or conducted under false pretenses. As a result, the user assumes that only vital information will be collected when signing up for or using a service, even while this information is used to gain knowledge of other users or inform other purposes that have not been initially declared.
    - **Social Pyramid** uses *Forced Communication or Disclosure* as a type of *Forced Action* to manipulate existing users into recruiting new users to use a service, often by tying this recruitment to additional functionality or other benefits. As a result, the user assumes that social recruiting is necessary to continue to use aspects of the service, even while this information is primarily used to build the service's user base.
  - **Gamification** subverts the user's expectation that system functionality is based on alignment with user goals and needs, instead coercing them into gaining access to aspects of a service through repeated (and perhaps undesired) use of aspects of the service.
    - **Pay-to-Play** uses *Gamification* as a type of *Forced Action* to initially claim that aspects of a service or product are available via purchase or download, but then later charging users to actually obtain that functionality. As a result, the user incorrectly assumes that a service or product will allow them certain functionality, leading to them downloading or purchasing the product or service under false pretenses.
    - **Grinding** uses *Gamification* as a type of *Forced Action* to require repeated, often cumbersome and labor-intensive actions over time in order to obtain certain relevant functionality. As a result, the user may seek to avoid these repetitive actions, leading to them making unwanted additional in-app purchases to unlock the same functionality without "grinding" over an extended period of time.
  - **Attention Capture** subverts the user's expectation that they have rational control over the time they spend using a system, instead tricking them into spending more time or other resources to continue use for longer than they otherwise would.
    - **Auto-Play** uses *Attention Capture* as a type of *Forced Action* to automatically play new video after an existing video has completed. As a result, the user may lose control over their viewing experience, leading them to watch more content than they intended or result in them watching content that is unexpected or harmful.
- **Social Engineering** is a strategy which presents options or information that causes a user to be more likely to perform a specific action based on their individual and/or social cognitive biases, thereby leveraging a user's desire to follow expected or imposed social norms.
  - **Scarcity or Popularity Claims** subverts the user's expectation that information provided about a product's availability or desirability is accurate , instead pressuring the user to purchase a product without additional reflection or verification.
    - **High Demand** uses *Scarcity and Popularity Claims* as a type of *Social Engineering* to indicate that a product is in high-demand or likely to sell out soon, even though that claim is misleading or false. As a result, the user may assume that demand is high when it is not, leading to their uninformed purchase of a product or service.
  - **Social Proof** subverts the user's expectation that the indicated behavior of others in a specific situation is correct or desirable, instead accelerating user decision-making and encouraging the user to trust flawed implications through provided information.
    - **Low Stock** uses *Social Proof* as a type of *Social Engineering* to indicate that a product is limited in quantity, even though that claim is misleading or false. As a result, the user may assume that a product is desirable due to demand, leading to undue or uninformed pressure to buy the product immediately.
    - **Endorsements and Testimonials** use *Social Proof* as a type of *Social Engineering* to indicate that a product or service has been endorsed by another consumer, even though the source of that endorsement or testimonial is biased, misleading, incomplete, or false. As a result, the user may assume that the endorsement or testimonial is accurate and unbiased, leading to their uninformed purchase of a product or service.
    - **Parasocial Pressure** uses *Social Proof* as a type of *Social Engineering* to indicate that a product or service has been endorsed by a celebrity, influencer, or other entity that the user trusts, even though the source of that endorsement is biased, misleading, incomplete, or false. As a result, the user may assume that the endorsement is accurate and unbiased, leading to their uninformed purchase of a product or service.
  - **Urgency** subverts the user's expectation that information provided about discounts or a limited-time deal for a product is accurate, instead accelerating the user's decision-making process by demanding immediate or timely action.
    - **Activity Messages** use *Urgency* as a type of *Social Engineering* to describe other user activity on the site or service, even though the data presented about other users' purchases, views, visits, or contributions are misleading or false. As a result, the user may

falsely feel a sense of urgency, assuming that others users are purchasing or otherwise interested product or service, leading to their uninformed purchase of a product or service.

  - **Countdown Timers** use *Urgency* as a type of *Social Engineering* to indicate that a deal or discount will expire by displaying a countdown clock or timer, even though the clock or timer is completely fake, disappears, or resets automatically. As a result, the user may feel undue urgency and purchasing pressure, leading to their uninformed purchase of a product or service.
  - **Limited Time Messages** use *Urgency* as a type of *Social Engineering* to indicate that a deal or discount will expire soon or be available only for a limited time, but without specifying a specific deadline. As a result, the user may feel undue urgency and purchasing pressure, leading to their uninformed purchase of a product or service.

- **Personalization** subverts the user's expectation that products or service features are offered to all users in similar ways, instead using personal data to shape elements of the user experience that manipulate the user's goals while hiding other alternatives.
  - **Confirmshaming** uses *Personalization* as a type of *Social Engineering* to frame a choice to opt-in or opt-out of a decision through emotional language or imagery that relies upon shame or guilt. As a result, the user may be convinced to change their goal due to the emotionally manipulative tactics, resulting in being steered away from making a choice that matched their initial goal.

B ANALYZED TAXONOMIES OF DARK PATTERNS

Table 2: Academic taxonomies of dark patterns.

	High-Level Pattern	Low-Level Pattern
Brignull 2018-2022 [6]	—	Sneak into Basket, Bait and Switch, Roach Motel, Price Comparison Prevention, Disguised Ads, Privacy Zuckering, Trick Questions, Hidden Costs, Confirmshaming, Friend Spam, Forced Continuity, Misdirection
Brignull 2023 [7]	—	Comparison Prevention, Confirmshaming, Disguised Ads, Fake Scarcity, Fake Social Proof, Fake Urgency, Forced Action, Hard to Cancel, Hidden Costs, Hidden Subscription, Nagging, Obstruction, Preselection, Sneaking, Trick Wording, Visual Interference
Bösch et al. [4]	Obscure	Privacy Zuckering, Immortal Accounts, Hidden Legalese Stipulations, Bad Defaults
	Maximize	Shadow User Profiles, Address Book Leeching, Forced Registration
	Deny	Immortal Accounts
	Preserve	Shadow User Profiles, Address Book Leeching
	Centralize	Shadow User Profiles
	Publish, Violate, Fake	—
Gray et al. [24]	Nagging	—
	Sneaking	Intermediate-Level Currency, Roach Motel, Price Comparison Prevention
	Obstruction	Bait and Switch, Sneak into Basket, Hidden Costs, Forced Continuity
	Interface Interference	Toying with Emotion, Aesthetic Manipulation, Trick Questions, Preselection, Disguised Ad, Hidden Information, False Hierarchy
	Forced Action	Gamification, Privacy Zuckering, Social Pyramid
Mathur et al. [40]	Sneaking	Sneak into Basket, Hidden Costs, Hidden Subscription
	Urgency	Limited-time Message, Countdown Timer
	Misdirection	Confirmshaming, Visual Interference, Trick Questions, Pressured Selling
	Social Proof	Activity Message, Testimonials
	Scarcity	Low-stock Message, High-demand Message
	Obstruction	Hard to Cancel
	Forced Action	Forced Enrollment
Luguri et al. [37]	Nagging	—
	Social Proof	Testimonials, Activity Messages
	Obstruction	Immortal Accounts, Intermediate-Level Currency, Roach Motel, Price Comparison Prevention
	Sneaking	Bait and Switch, Sneak into Basket, Hidden Costs, Hidden Subscription / Forced Continuity
	Interface Interference	Cuteness, False Hierarchy / Pressured Selling, Toying with Emotion, Trick Questions, Preselection, Disguised Ad, Hidden Information / Aesthetic Manipulation, Confirmshaming
	Forced Action	Friend spam/social pyramid/address book leeching, Privacy Zuckering, Gamification, Forced Registration
	Scarcity	High Demand Message, Low Stock Message
	Urgency	Countdown Timer, Limited Time Message

Table 3: Regulatory taxonomies of dark patterns.

	High-Level Pattern	Low-Level Pattern
EDPB [17]	Overloading	Continuous Prompting, Privacy Maze, Too Many Options
	Skipping	Deceptive Snugness, Look Over There
	Stirring	Emotional Steering, Hidden in Plain Sight
	Obstructing	Dead End, Longer than Necessary, Misleading Action
	Fickle	Lacking Hierarchy, Decontextualizing, Language Discontinuity, Inconsistent Interface
	Left in the Dark	Conflicting Information, Ambiguous Wording or Information
EU Com. (EC) [15]	Nagging	—
	Social Proof	Testimonials, Activity Messages
	Obstruction	Intermediate-Level Currency, Roach Motel / Difficult Cancellations, Price Comparison Prevention
	Sneaking	Bait and Switch, Sneak into Basket, Hidden Costs, Hidden Subscription / Forced Continuity
	Interface Interference	Toying with Emotion, Trick Questions, Preselection (default), Disguised Ad, Hidden Information / False Hierarchy, Confirmshaming
	Forced Action	Forced Registration
	Urgency	Countdown Timer / Limited TIme Message, Low Stock / High Demand Message
OECD [48]	Forced Action	Forced Registration, Forced Disclosure / Privacy Zuckering, Friend Spam / Social Pyramid / Address Book Leeching, Gamification
	Interface Interference	Hidden Information, False Hierarchy, Preselection, Misleading Reference Pricing, Trick Questions, Disguised Ads, Confirmshaming / Toying with Emotion
	Nagging	Nagging
	Obstruction	Hard to Cancel or Opt Out / Roach Motel / Click Fatigue / Ease, (Price) Comparison Prevention, Immortal Accounts, Intermediate Currency
	Sneaking	Sneak into Basket, Hidden Costs / Drip Pricing, Hidden Subscription / Forced Continuity, Bait and Switch (including Bait Pricing)
	Social Proof	Activity Messages, Testimonials
	Urgency	Low Stock / High Demand Message, Countdown Timer / Limited Time Message
UK CMA [10]	Choice Structure	Defaults, Ranking, Partitioned Pricing, Sludge, Bundling, Dark nudge, Choice overload and decoys, Virtual currencies in gaming, Sensory manipulation, Forced outcomes
	Choice Information	Drip pricing, Reference pricing, Framing, Complex language, Information overload
	Choice Pressure	Scarcity and popularity claims, Prompts and reminders, Messengers, Commitment, Feedback, Personalisation
US FTC [21]	Endorsements (Social Proof)	False Activity Messages, Deceptive Consumer Testimonials, Deceptive Celebrity Endorsements, Parasocial Relationship Pressure
	Scarcity	False Low Stock Message, False High Demand Message
	Urgency	False Discount Claims, False Limited Time Message, Baseless Countdown Timer
	Obstruction	Immortal Accounts Roadblocks to Cancellation, Price Comparison Prevention
	Sneaking or Information Hiding	Intermediate Currency, Hidden Subscription or Forced Continuity, Drip Pricing, Hidden Costs, Hidden Information, Sneak-into-Basket
	Interface Interference	Bait and Switch, Disguised Ads, False Hierarchy or Pressured Upselling, Misdirection
	Coerced Action	Friend Spam, Social Pyramid Schemes, and Address Book Leeching, Pay-to-Play or Grinding, Forced Registration or Enrollment, Nagging, Auto-Play, Unauthorized Transactions
	Asymmetric Choice	Subverting Privacy Preferences, Preselection, Confirm Shaming, Trick Questions