

Understanding Dark Patterns in Home IoT Devices

Monica Kowalczyk*
kowalczyk.m@northeastern.edu
Northeastern University
Boston, Massachusetts, USA

Johanna Gunawan*
gunawan.jo@northeastern.edu
Northeastern University
Boston, Massachusetts, USA

David Choffnes
choffnes@ccs.neu.edu
Northeastern University
Boston, Massachusetts, USA

Daniel J. Dubois
d.dubois@northeastern.edu
Northeastern University
Boston, Massachusetts, USA

Woodrow Hartzog
w.hartzog@northeastern.edu
Boston University
Boston, Massachusetts, USA

Christo Wilson
cbw@ccs.neu.edu
Northeastern University
Boston, Massachusetts, USA

ABSTRACT

Internet-of-Things (IoT) devices are ubiquitous, but little attention has been paid to how they may incorporate dark patterns despite consumer protections and privacy concerns arising from their unique access to intimate spaces and always-on capabilities. This paper conducts a systematic investigation of dark patterns in 57 popular, diverse smart home devices. We update manual interaction and annotation methods for the IoT context, then analyze dark pattern frequency across device types, manufacturers, and interaction modalities. We find that dark patterns are pervasive in IoT experiences, but manifest in diverse ways across device traits. Speakers, doorbells, and camera devices contain the most dark patterns, with manufacturers of such devices (Amazon and Google) having the most dark patterns compared to other vendors. We investigate how this distribution impacts the potential for consumer exposure to dark patterns, discuss broader implications for key stakeholders like designers and regulators, and identify opportunities for future dark patterns study.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI); Interaction design;** • **Security and privacy** → Human and societal aspects of security and privacy; • **Social and professional topics** → *Computing / technology policy.*

KEYWORDS

UX design, dark patterns, IoT, human factors

ACM Reference Format:

Monica Kowalczyk, Johanna Gunawan, David Choffnes, Daniel J. Dubois, Woodrow Hartzog, and Christo Wilson. 2023. Understanding Dark Patterns in Home IoT Devices. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23), April 23–28, 2023, Hamburg, Germany*. ACM, New York, NY, USA, 27 pages. <https://doi.org/10.1145/3544548.3581432>

*Both authors contributed equally to this research.

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 '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9421-5/23/04...\$15.00

<https://doi.org/10.1145/3544548.3581432>

1 INTRODUCTION

Internet-of-Things (IoT) devices have become ubiquitous, offering a wide range of functionality including home automation, voice assistance, media playback, video surveillance, appliances, and health monitoring. Despite extensive prior work on the security [11, 32, 57, 86, 87, 98, 100] and privacy [14, 22, 23, 27, 62–65, 85] implications of these purpose-built hardware devices, little attention has been paid to how the the unique experiences such devices present may incorporate potentially deceptive and harmful designs.

Such designs are often called “dark patterns,” which interfere with user behavior to capture people’s attention, extract people’s consent to boilerplate contracts, goad people into financial transactions, and nudge people into exposing or sharing personal information, among myriad other unintended or negative outcomes. Scholars have systematically identified and measured dark patterns in primarily visual modalities like apps and websites, demonstrating their wide prevalence and potential harms to user privacy, autonomy, finances, and cognitive resources. These studies provide taxonomies of dark patterns that align patterns to different traits, e.g., user outcome attributes [68], design approaches [26, 39], and interaction context [42]. With access to sensitive information (e.g., health data, video feeds, sensor data), always-on capabilities, and experiences that span hardware and software, IoT device experiences may exacerbate harms or include previously undocumented dark pattern instances. For example, the *detour* dark patterns observed in this study unexpectedly direct users outside of the immediate modality (e.g., companion app or on-device interface), which interrupts or obstructs the user’s intended behavior and can risk privacy and autonomy harms.

In this paper, we conduct a systematic study of IoT-device dark patterns across multiple device types. We investigate a diverse set of 57 popular IoT devices spanning nine categories, three interaction methods (app, voice, and direct device control), and six manufacturers. IoT devices offer unique lenses through which to study dark patterns because the devices offer a wide range of functionality, cut across diverse real-world contexts, have unchangeable physical-world interfaces (e.g., buttons, voice, and screens), and intersect with app-based modalities when companion apps are required. Our study seeks to answer the following questions:

- (1) *To what extent do dark patterns observed in other modalities apply to IoT? Do IoT modalities change our understanding of dark patterns and give rise to new dark patterns?* IoT devices offer new interaction methods (e.g., physical and primarily

voice-controlled interfaces) compared to websites and apps. This potentially makes them subject to existing dark patterns, as well as new ones that depend on the new modalities.

- (2) *How do dark patterns in IoT devices relate to device type, manufacturer, modality, and the context in which interactions were performed?* IoT devices exhibit diversity along several dimensions, including interaction methods and offered functionality. Understanding the relationship between such diversity and dark patterns can help predict the presence of dark patterns based on IoT device characteristics.
- (3) *What are the implications of observed IoT dark patterns?* Given the observed IoT dark patterns, we posit why certain dark patterns manifest in the IoT environment, and note key challenges and identify trends that can guide further investigation into IoT dark patterns.

To answer these questions, we address key methodological challenges for studying dark patterns in the IoT environment, including how to select a reasonable set of devices to test from the wide range of device types and manufacturers on the market, how to conduct rigorous testing across such devices with a tractable amount of effort and time, and how to collect reproducible and reusable datasets from experiments with these devices. We build a codebook that expands upon prior taxonomies [26, 42] with 12 new patterns that we discovered during our pilot experiments, organize our codebook by interaction context to facilitate human annotation and capture which contexts carry higher dark pattern risk, and tackle strategies for labeling dark pattern co-occurrences. This results in a set of scripted interactions with IoT devices recorded on video, with annotations that include timestamps when dark patterns occurred and their position in a video frame.

We present an analysis of the dark patterns identified in our experiments. On average we found 10–11 unique dark patterns per device, and noted between 3–90 total patterns (accounting for multiple instances of patterns) per device. We find that pre-selection and visual preference patterns are most common adopted, as well as patterns related to account creation, consent/permissions, and account deletion. We find crdiffmore dark patterns in always-on devices like cameras, doorbells, and speakers than in other types. We find that Amazon and Google devices tend to contain more dark patterns. Our results do not point to one single factor as the primary driver of dark patterns; rather, our findings highlight the necessity of multi-factor analysis for dark patterns. We also discuss how nontransparent designs in IoT devices may exacerbate financial and privacy harms, as well as other risks from the IoT context.

Finally, we discuss the implications of our findings, focusing on key aspects of the IoT environment that give rise to observed behaviors, challenges for future research on IoT dark patterns, and interesting trends identified in our dataset. We suggest potential mitigations for design practitioners and regulators, noting strategies to minimize darkness and improve transparency.

2 BACKGROUND

We begin by reviewing related work on dark patterns in general, and IoT risks, harms, and user experience (UX) in particular. We contextualize and motivate this study within this broader scope of existing scholarship.

2.1 Dark Patterns

Dark patterns [20] are user interface designs that trick users into unwanted or unintentional behavior, typically against users' best interests. Conceptually, dark patterns relate to malicious interfaces [25], online manipulation [93], nudges [94], and UX design [47]. Dark patterns have received public attention in the press [58, 73, 84], scholarly and regulatory workshops [21, 59, 79], and government reports [24, 33, 34]. Commensurate with this increasing awareness, dark patterns are now regulated in some contexts such as consent interactions [4, 6, 51].

Academics have developed robust taxonomies of dark patterns based on their underlying mechanisms or tactics, both from a design perspective [39, 67] and with a privacy lens [19]. Following these taxonomies, observational and measurement studies identify and enumerate dark patterns in app [26, 37, 39, 42] and website [39, 42, 67] modalities. These studies demonstrate dark pattern pervasiveness and the diversity of designs across various user interactions, platforms, and modalities. Other work delves into specific contexts, providing detailed insight into certain dark pattern or interaction types. These contexts include e-commerce [67], consent interactions and cookie banners [40, 41, 44, 45, 52, 61, 78, 92], account deletion interactions [56, 89], online addictions [3, 72], and specific online services [56, 71]. The level of detail explored in these studies provides evidence for regulatory responses. For example, Gray et al. [40] analyzed dark pattern use in the non-compliance of consent management providers (CMPs) to the GDPR consent requirements; such evidence can affirm consumer complaints in enforcement actions [1, 76]. A key open question, which we discuss below, is the extent to which prior work on other modalities and contexts apply to the IoT environment.

Scholars taxonomized the range of poor outcomes and consumer harms dark patterns may cause [43, 47, 68, 75, 91]. User studies capture consumer reactions to dark patterns [18, 26, 60, 80, 89] and find that people do feel manipulated or disadvantaged by dark patterns [38, 71], but that people vary in how they perceive dark patterns and their theorized harms (e.g., some are surprised by certain dark patterns, while others are unsurprised but resigned).

Researchers are beginning to grapple with the role that *context* plays with respect to dark patterns. For example, work focusing on *Roach Motels* [20, 39] (i.e., designs that make it easy to get into a situation such as a subscription, but hard to get out of) frame dark patterns as socio-technological phenomena, largely dependent on how they are interacted with rather than how they are presented [16]. Gray et al. [40] stress an “n-dimensional” approach for researching dark patterns that incorporates factors like time, interaction, design, psychology, and law via multi-disciplinary analysis. Work by Gunawan et al. [42] embraces this approach by comparing dark patterns across thematic UX categories, while work by Mathur et al. [68] critiques dark patterns through a variety of disciplinary lenses. Our study aims to continue this line of scholarship by investigating dark patterns in previously unexplored contexts and interaction modalities with multi-factor analyses.

2.2 IoT Contexts and Emerging Modalities

Context is critical when considering dark patterns in IoT devices that may serve vastly different purposes. They can access intimate

spaces (e.g., bedrooms or bathrooms), collect highly specific data through sensors, and be perpetually on or listening (e.g., microphones in voice assistants), with device hardware that may not provide feedback to users (e.g., not indicating when a device is recording video). These qualities present unique and additional challenges for user privacy, security, and safety (as demonstrated by extensive scholarship on IoT privacy and security issues [7, 9, 10, 12, 13, 17, 23, 27, 30, 31, 49, 53, 62–65, 69, 69, 74, 81, 85, 88, 95, 97, 99, 101]) relative to websites and apps.

Dark designs may exacerbate these safety and security concerns. Limited surfaces for interaction on IoT devices, or complex control schemes involving companion apps, may offer opportunities for designers to obfuscate, discourage the use of, or even omit privacy-critical functionality. Recent work has begun to measure users' interactions and perceptions of settings interfaces in IoT devices [15, 66], providing early indicators of "good" UX in IoT devices. Owens et al. [80] investigated non-visual interfaces like voice assistants through speculative design fiction exercises and user surveys, identifying that while participants found intentionally deceptive scenarios to be more problematic than those that were not, overall the participants did not show much concern over deceptive scenarios.

IoT devices may use the same platform as mobile phones (e.g., using Android variants such as Android Auto, Android TV, and Android Wear). An open question, then, is whether such devices exhibit similar dark patterns as those found in mobile apps that run on the same OS foundation [26, 42]. While this might be the case for some devices, several factors may alter the frequency or types of dark patterns adopted in IoT devices when compared to mobile apps. For example, when an IoT device utilizes a commodity OS like Android, the user interface and interaction modalities may differ from smartphone Android (e.g., a remote for Android TV rather than a touchscreen). Furthermore, many IoT devices run bespoke or uncommon OSes (e.g., Tizen from Samsung or Fuschia from Google), or include sensors that are not present on smartphones (e.g., always-on cameras).

2.3 Building On Prior Dark Patterns Work

Our work lies at the intersection of dark patterns and IoT. Our codebook for identifying dark patterns (see subsection 3.3) draws on taxonomies and harm frameworks from prior work [19, 42, 67, 68], and contributes to early explorations on darkness in emerging modalities [80].

Like prior measurement work, we manually interact with devices, look for dark patterns, and link dark patterns to potential harms. We draw on approaches [26, 42] and design perspectives [15] from prior work to inform our process for interacting with IoT devices and labeling dark patterns (see subsection 3.2). We expand upon existing measurements by collecting frequencies of repeated dark pattern encounters. Further, we directly compare our IoT dark pattern measurements to prior manual studies [26, 42] (see subsection 4.1.4).

Previous work primarily focused on visual modalities like apps and websites [26, 42, 67]. In contrast, we holistically examine device experiences multimodally, through direct interaction with the devices, voice interactions, and interactions with companion

apps. Prior manual studies conducted uniform time-bound actions per app or service [26, 42], which were naturally constrained by each modality's affordances (e.g., all observed apps or mobile sites were interacted with via touchscreen [26, 42] or desktop sites were viewed on a computer). To understand a device' experience across offered modalities and robustly explore available features or necessary configuration, our methods necessitated unrestricted interaction durations and flexible interaction scripts. As such, we do not conduct disparate, per-modality interactions nor draw apples-to-apples comparisons between the modalities inspected. Specific to smart devices, we depart from recent design fiction and user study [80] approaches for exploring dark patterns in voice interface, instead using manual testing methods on real device interfaces and examining modalities beyond voice alone.

3 METHODS

We now describe the methods used in our study. This section covers preliminary experiments, then describes how we arrived at the final number of devices examined, how we inspected and interacted with each device, and how we annotated dark patterns per IoT modality. We also present validation of our annotation process.

3.1 Lab Environment and Devices

Our study was conducted on devices purchased between 2017–2022 in the United States and primarily housed within a single, controlled-access lab environment. This lab environment resembles a studio apartment, with devices installed in the manner they might be in a typical home. Two TV devices were hosted off-site by a trusted third-party. Including these two TVs, we studied 57 devices spanning nine broad types: *home automation*, *home appliance*, *health*, *smart hub*, *camera*, *doorbell*, *television*, *media device*, and *speaker* devices as listed in Table 1. Although additional devices were available in our lab environment, we excluded devices that were model-year iterations of the same device (e.g., of our available Echo Dots, we included only the most recent 4th Gen Echo Dot) and devices with dysfunctional factory-reset capability. We list excluded devices in the Appendix (Table 6).

In general, IoT devices offer multiple interaction modalities, including on-device buttons, touch screens, remotes, voice commands, or even no on-device interface at all. For devices in the last category, the only available modality was a companion app installed on a smartphone. To operate these devices, we primarily used Android phones with relevant companion apps installed, with exceptions for HomePods, which we paired with an iPhone.

3.2 Pilot Experiment

Two key challenges for identifying dark patterns in IoT devices are:

- (1) Existing *interactions scripts* that stipulate how to manually exercise and record interactions with websites and apps [26, 42] are unlikely to be sufficient for exercising the full functionality of IoT devices, given differences in modality affordances.
- (2) Existing codebooks for annotating dark patterns [26, 39, 42] may not be sufficiently illustrative when applied to IoT or multimodal device experiences.

To address these issues we conducted an exploratory pilot experiment to identify any necessary adjustments to prior methods for

Device Type	Device Name	Ecosystem	App Name (If Used)	App Dependency	Video Duration
Home Automation	Amazon Smart Plug	Amazon	Amazon Alexa	All interactions	0:11:13
	Jinwoo Smart Bulb		Jinwoo Smart	All interactions	0:19:22
	Gosund Smart Light Bulb		Gosund	All interactions	0:12:30
	Govee LED Light Bulb		Govee Home	All interactions	0:15:41
	Magichome Strip		Magic Home Pro	All interactions	0:09:24
	Meross Door Opener		meross	All interactions	0:12:17
	Nest Thermostat*	Google	Nest	All interactions	0:26:27
	Ring Chime	Amazon	Ring	All interactions	0:19:50
	Smartlife LED Bulb	Smartlife	Smart Life	All interactions	0:18:27
	WeMo Plug		Wemo	All interactions	0:13:14
	Thermopro TP90		ThermoPro Home	Smart interactions	0:06:59
	TP-Link Bulb	Kasa	Kasa Smart	All interactions	0:14:03
	TP-Link Plug	Kasa	Kasa Smart	All interactions	0:14:23
Camera	Amcrest Cam		Amcrest View Pro	All interactions	0:11:45
	Arlo Q Cam	Arlo	Arlo Secure: Home Security	All interactions	0:19:24
	D-Link Cam		mydlink	All interactions	0:13:15
	Lefun Cam		MIPC	All interactions	0:11:31
	Nest Camera	Google	Google Home	All interactions	0:13:13
	Ring Camera	Amazon	Ring	All interactions	0:22:42
	Ring Camera (Indoor)	Amazon	Ring	All interactions	0:16:04
	Tuya Smart Camera		Tuya Smart	All interactions	0:17:00
	Wyze Cam		Wyze	All interactions	0:23:47
	Yi Home Camera		Yi Home	All interactions	0:11:08
Media Device	Apple TV*	Apple		No interactions	0:17:33
	Chromecast w/ Google TV*	Google		No interactions	0:31:36
	Facebook Portal Mini*			No interactions	0:34:33
	Fire TV*	Amazon		No interactions	0:51:49
	Nintendo Switch*			No interactions	0:38:28
	Roku TV*			No interactions	0:25:20
TiVo Stream*			No interactions	0:39:37	
Smart Hub	Aqara Hub		Aqara Home	All interactions	0:19:37
	Sengled Smart Hub		Sengled Home	All interactions	0:12:12
	SmartThings Hub	Samsung	SmartThings	All interactions	0:34:43
	Switchbot Hub		SwitchBot	All interactions	0:10:47
	Philips Hue Bridge		Philips Hue	All interactions	0:13:23
Doorbell	Arlo Doorbell	Arlo	Arlo Secure: Home Security	All interactions	0:13:19
	Nest Doorbell	Google	Google Home	All interactions	0:16:33
	Ring Doorbell	Amazon	Ring	All interactions	0:15:21
	Ring Doorbell ('21, Wired)	Amazon	Ring	All interactions	0:28:07
Speaker	Echo Dot (4th Gen)	Amazon	Amazon Alexa†	Setup interactions only	0:48:00
	Echo Show 5*	Amazon	Amazon Alexa†	Smart interactions	0:44:09
	Home Mini	Google	Google Home†	Setup interactions only	0:18:37
	Nest Mini	Google	Google Home†	Setup interactions only	0:45:15
	Homepod	Apple	Home (iPhone)†	Setup interactions only	0:21:56
	Homepod Mini	Apple	Home (iPhone)†	Setup interactions only	0:31:12
	Nest Hub Max*	Google	Google Home†	Setup interactions only	0:41:19
Home Appliance	Samsung Fridge*	Samsung		No interactions	0:46:25
	GE Microwave		SmartHQ	Smart interactions	0:20:33
TV	LG TV*			No interactions	0:27:51
	Samsung TV*	Samsung		No interactions	0:22:42
	Sony TV*	Sony		No interactions	0:30:00
	Vizio TV*	Vizio		No interactions	0:21:49
Health	Oxylink Oxygen Monitor		ViHealth	All interactions	0:12:27
	Renpho Smart Scale		Renpho	Smart interactions	0:18:54
	Withings BPM Connect	Withings	Withings Health Mate	All interactions	0:24:49
	Withings Sleep	Withings	Withings Health Mate	All interactions	0:09:15
	Withings Thermo	Withings	Thermo	All interactions	0:10:20

Table 1: The 57 devices used in this study. Device names marked with asterisks (*) contained navigable screens in the device hardware. App names marked with daggers (†) denote devices for which we annotated and discovered dark patterns in both the device and the app. We collected approximately 20 hours of recordings in total. Refer to Table 5 in the Appendix A for device firmware or app software information.

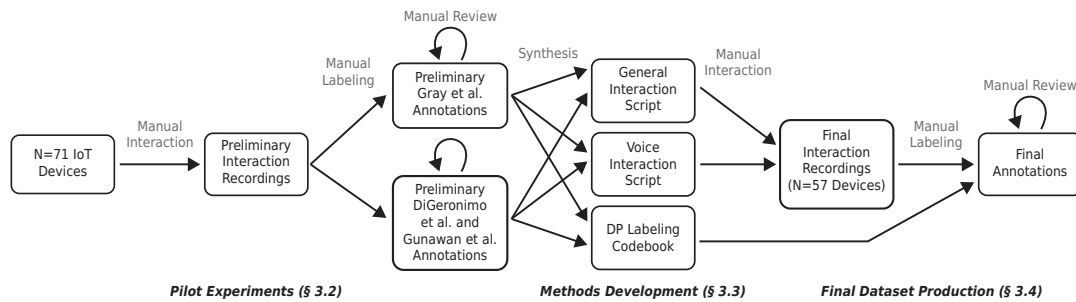


Figure 1: Flowchart detailing our manual analysis procedures, from pilot experiments to methods adjustments and final annotations. Note that we added 13 more devices a few months after performing the preliminary interaction recording, bring the total from N=71 to N=84 devices.

our final set of experiments, as well as identify novel dark patterns for our codebook. Figure 1 presents a flowchart of our pilot experiments.

The first author interacted with all devices housed in the lab during August–October 2021 (N=71), using a Google Pixel 3 phone or iPhone 6 to control devices with an available companion app. Screen recordings were taken for companion app interactions, and video recordings were taken for physical device interactions. All phones and devices were connected to the local lab network while each device was examined, and were logged-in to pre-existing lab user accounts per-device. These user accounts varied in age as they were created as needed over several years. As such, many devices contained existing usage data. Each device was used as intended (e.g., the fridge was used to store food, light bulbs were turned on and off, cameras took or displayed footage) and we explored all navigation options provided by a device (e.g., menus in visual interfaces like apps or physical touchscreens, and queries in voice interfaces). For each device we sought to examine every navigable feature or area in the interface, and followed navigation to the depth of any end-node pages or frames. We did not attempt to interact with every possible button or toggle provided by end pages or frames (either on the device itself or in its companion app).

The first author then reviewed each video (either a screen recording or external footage) and labeled dark patterns in two iterations. First, we labeled unique dark patterns according to the five main categories in the Gray et al. [39] taxonomy and kept written descriptions of encountered patterns. These labels were discussed with two co-authors to achieve alignment on dark pattern identification. Second, the first author labeled each video according to a codebook based on specific dark pattern instances from Di Geronimo et al. [26] and Gunawan et al. [42], noting dark behaviors that had not been explicitly described in prior work. These notes were reviewed and discussed with the second author and resulted in the addition of new labels for novel dark pattern cases.

As manual annotations are subject to individual bias, the first two authors independently annotated a subset of these device recordings (at least one device per type, plus an additional speaker with a touchscreen; N=8 devices). The two authors then compared labels for validation. Disagreements were discussed and corrected towards an agreed-upon understanding of each pattern, generally defaulting to the first author’s labels.

3.3 Methods Development

Next we refined our interaction scripts and codebooks for the final round of device interactions, data collection, and subsequent annotations. Figure 1 shows the output of this process: a general interaction script used for non-voice controlled IoT devices and companion apps, an interaction script for voice-controlled devices, and an expanded codebook for labeling dark patterns.

3.3.1 Environment Isolation. Compared to a neatly isolated environment for testing websites and apps, the single-network, live lab environment presented unique challenges for IoT data collection. Pilot study interactions revealed the potential influence of interaction history or pre-existing data, and logged-in devices prevented insight into device and account setup experiences. With all devices on the same network, some devices and apps could communicate with each other (e.g., light bulbs connected to several smart hub apps, or Amazon devices connected to the same app). This blurred distinctions between similar devices: in some cases, we received notifications from devices we were not intending to interact with if that device shared the same app as the device we were inspecting.

To mitigate these issues during our final data collection, we factory reset each device and—like prior work [26, 42]—created a fresh user account for each device as needed.¹ We also provisioned a separate, isolated network solely for this study, such that only one controller phone and the currently-examined device would be connected to the network at the same time (all other lab devices remained on the original network).

3.3.2 Embedded Browsers and OS Interfaces. Some devices and companion apps loaded web pages using built-in browsers. Unlike prior work [42], we included dark patterns discovered in such web pages if these (1) transmitted login information or (2) delivered information promised by previous device/app menus or features, but did not count dark patterns in embedded pages that served tertiary purposes. For example, the Google Nest Mini app contains menu items that lead to a logged-in shopping web page hosted by Google that allows users to purchase other Google Home products. In this case, we inspected web pages relevant to the shopping task, but did not visit Google Account pages or other product pages hyperlinked from the same page. This restriction avoids incorporating the

¹For Google or iPhone devices using OS-level apps like Google Home or Apple Home, fresh user accounts were used for smartphone login.

entirety of vendor websites into our analysis of specific IoT products. Similarly, some devices and controlling smartphones share the same manufacturer (e.g., HomePod and iPhone; Google Home with Android devices) and leverage both OS-level and companion-app-level interfaces to control the smart device. Thus, we retained dark patterns found in the OS in our dataset if they were encountered naturally within our interaction script and did not explore other OS interfaces beyond those relevant to IoT interactions.

3.3.3 Speaker Interactions. In the pilot study, we asked our speakers a small set of common queries, suggested verbatim by device documentation and lists from popular consumer electronics websites [36, 50, 83]. These lists, however, tended to mention highly specific one-off use cases like shopping for an item, playing music, or inquiring about the weather. Other desired interactions, like navigating settings or account deletion, were not as readily found.

To keep our speaker interactions as consistent as possible with physical devices and companion apps, we constructed our own list of commands that would exercise the sample commands listed above, as well as attempt to navigate settings and conduct exit interactions (e.g., delete the associated account). We conducted smaller tests of these commands across all three speaker ecosystems in our study to identify ecosystem-specific functionality (like the presence of a guest mode, or the ability to mute the mic by voice command) and test the limits of available voice-controlled actions. Account and device setup interactions were omitted from this list, as we discovered that all speakers in our tests required companion-app-based setup prior to use.

3.3.4 Novel Dark Patterns. Based on the results of the pilot experiment we added 12 new dark patterns to our codebook. Throughout the pilot experiments, the authors noted all unlabeled device and companion-app behaviors that fit within one or more broad classifications of dark patterns [39, 67]. The authors discussed these behaviors while co-annotating recordings, comparing independent annotations, and when reviewing intermediary analyses. The authors agreed that these behaviors fell within high-level dark pattern categorizations and traits, but were not specifically captured by prior codebooks used during the pilot experiment [26, 42]. We present novel dark patterns in **bold** in Table 2. We discuss these novel patterns and their implications in subsection 4.2.

3.4 Final Dataset Production

We now describe our final device interaction, interaction recording, and video annotation procedure, which took place October 2021–June 2022. We examined the 57 devices shown in Table 1.

3.4.1 Navigating Interfaces. Our approach to interactions relied on interaction scripts drawn from prior work [26, 42] with modifications stated in subsection 3.3. Our script was designed to uncover and explore as many possible features—including settings categories—afforded by each smart device across companion apps and device hardware (which included voice-controlled and visual interfaces). We conducted device and account setup, traversed available features and settings, and performed exit interactions (e.g., logout, device disassociation, data or account deletion, etc.) where possible for each device. When setting up an account we agreed to all options that were preselected or preferred in visual hierarchies.

In cases where dark patterns did not steer us towards particular choices we chose the first available option, from top to bottom and from left to right. We focused on traversing as many main features or options as were intuitively provided within the companion app or device interface. Likewise, we conducted a best-effort approach to visit all available settings, subject to limitations where a device’s settings navigation would require an unusually long time to traverse (e.g., each setting was individually paginated, requiring multiple page loads to traverse). In these cases we visited a subset of settings for reasonable coverage.

3.4.2 Device Interactions. We factory-reset each device, connected them to the isolated network, configured devices using a unique e-mail address for registration if required, and used factory-reset Google Pixel 3 and iPhone 6 phones for any app-based interactions. We recorded companion app interactions using screen recording software and took video footage of device interactions using a smartphone camera on a tripod.

Each IoT device test began with an attempt to interact with the physical device. When prompted by the device or when apparent that the device required a companion app for smart features, we installed the relevant app on a compatible smartphone, began screen recording, and interacted with the device through the app. We followed device or app guidance when determining which modality to use and traversed any remaining features or device-specific settings in the device or app after completing as many available actions as possible for the primary purpose(s) of the device. Following our traversal, we examined any available app-level or physical-device settings and attempted any available exit interactions.

3.4.3 Companion App Interactions. To preserve ecological validity, we sought to interact with IoT devices as directly as possible, as an average user might. Thus we refrained from using companion apps unless it was the only way to control a device or the device required us to. Of our 57 devices, only 12 could be fully interacted without a companion app: the media devices (e.g., Apple TV and Roku TV), the TVs, and the fridge. If a device prompted login but did not require the companion app, we used a desktop browser for account registration.

We identified cases where the same companion app controlled multiple devices in our tests, which allowed us to save time without loss of coverage by interacting in full only once for all its corresponding devices. For example, both Alexa speakers and the Amazon Plug use the *Amazon Alexa* app; the Ring Camera, Ring Chime, and Ring Doorbell all share the *Ring* app; all three Google speakers and the Chromecast use *Google Home*, etc. In these cases we fully traversed each app only once, and otherwise only interacted with the app as necessary per-device, on demand—typically for fresh account or device setup, or managing relevant settings.

3.4.4 Annotation Procedure and Validation. We manually annotated the video recordings produced by our device and companion app interactions for dark patterns using the codebook in Table 2. Prior studies using similar methods operationalized dark patterns as binary variables that were either present or not present in each sample [26, 42]. In contrast, Mathur et al. [67] counted the number of each type of dark pattern that appeared on each website and web page during automated crawls. We use both approaches in

Context Category	Dark Pattern Description	Mapping to Prior Taxonomies	Potential Harms
Registration	Account required to use service	Forced Registration [19], Forced Action [26, 39]	Privacy [68]
	Account required to set up device	Forced Registration [19], Forced Action [26, 39]	Privacy [68]
Engagement	Gamification	Gamification [26, 39]	Cognitive [68]
	Extraneous notification badges Extraneous message centers Extraneous social media features	Aesthetic Manipulation [26, 39] Nagging [26, 39] Nagging [26, 39]	Cognitive [68] Cognitive [68] Cognitive [68]
Consent and Permissions	No Terms of Service/Privacy Policy	Privacy Zuckering [19, 20, 26, 39]	Privacy, Autonomy [68]
	No link to Terms of Service/Privacy Policy	Hidden Legalese Stipulations [19], Hidden Information [26, 39]	Privacy, Autonomy [68]
	No consent checkbox for Terms of Service/Privacy Policy	Privacy Zuckering [19, 20, 26, 39]	Privacy, Autonomy [68]
	Consent checkbox is preselected	Bad Defaults [19], Preselection [26, 39]	Privacy, Autonomy [68]
	Consent notice includes email subscription	Bad Defaults [19], Preselection [26, 39]	Autonomy [68]
Ads	Preselected email subscription checkbox	Bad Defaults [19], Preselection [26, 39]	Autonomy [68]
	Permission requested without explanation	Hidden Legalese Stipulations [19], Hidden Information [26, 39]	Privacy, Autonomy [68]
	Permission pops up unprompted	Nagging [26, 39]	Cognitive, Privacy [68]
	Device sensed without permissions	Privacy Zuckering [19, 20, 26, 39]	Privacy, Autonomy [68]
	Nonpermanent opt out	Trick Question [20, 26, 39]	Autonomy [68]
Money	Native ads	Disguised Ads [20, 26, 39]	Cognitive [68]
	Hard to close ads	Aesthetic Manipulation [26, 39]	Cognitive [68]
	Inconsistent close buttons	Aesthetic Manipulation [26, 39]	Cognitive [68]
	Interact with ads to unlock a feature	Forced Action [26, 39]	Cognitive, Autonomy [68]
Shopping	Pay to avoid ads	Hidden Information [26, 39]	Financial [68]
	Pay for fictional currency	Intermediate Currency [26, 39]	Financial [68]
	Pay for badges	Intermediate Currency [26, 39]	Financial [68]
	Unsolicited free trial	Forced Continuity [20, 26, 39]	Autonomy [68]
	Free trial requires payment method	Forced Continuity [20, 26, 39]	Financial [68]
Seen in Settings	Pay for long term use	Forced Continuity [20, 26, 39]	Financial [68]
	Feature seems free but is not	Disguised Ads [20, 26, 39]	Financial [68]
	Feature seems premium but is not	Hidden Information [26, 39]	Financial [68]
	Cannot sort free from premium content	Aesthetic Manipulation [26, 39]	Cognitive [68]
	Suggests preferred items	False Hierarchy [26, 39]	Autonomy [68]
Leaving	Sneaking items into basket	Sneak Into Basket [26, 39]	Financial, Autonomy [68]
	Optional items are preselected	Sneaking [26, 39]	Financial, Autonomy [68]
	Shaming language when opting out	Privacy Zuckering [19], Toying with Emotion [26, 39]	Autonomy [68]
	Item has a different price	Bait and Switch [20, 26, 39]	Financial [68]
	Surprise fees	Hidden Information [26, 39]	Financial [68]
Interface Interference	Countdown timer	Toying with Emotion [26, 39]	Financial, Autonomy [68]
	Social proof	Toying with Emotion [26, 39]	Financial, Autonomy [68]
	No bulk options for settings	Privacy Zuckering [19], Aesthetic Manipulation [26, 39]	Cognitive [68]
	No notification settings	Bad Defaults [19], Forced Action [26, 39]	Cognitive, Autonomy [68]
	No privacy settings	Bad Defaults [19], Forced Action [26, 39]	Privacy, Autonomy [68]
Subverting Expectations	Notification settings preselected	Bad Defaults [19], Preselection [26, 39]	Cognitive, Autonomy [68]
	Privacy settings preselected	Bad Defaults [19], Preselection [26, 39]	Privacy, Autonomy [68]
	Hard to navigate settings	Privacy Zuckering [19], Aesthetic Manipulation [26, 39]	Cognitive [68]
	Inconsistent Settings UI	Privacy Zuckering [19], Aesthetic Manipulation [26, 39]	Cognitive [68]
	Settings detour to a different modality	Privacy Zuckering [19], Forced Action [26, 39]	Cognitive [68]
Subverting Expectations	No logout	Immortal Accounts [19], Roach Motel [20, 26, 39]	Autonomy [68]
	No account deletion	Immortal Accounts [19], Roach Motel [20, 26, 39]	Privacy, Autonomy [68]
	Unclear deletion options	Privacy Zuckering [19], Roach Motel [20, 26, 39]	Privacy, Autonomy [68]
	Time delayed deletion	Immortal Accounts [19], Roach Motel [20, 26, 39]	Privacy [68]
	Cannot remove device	Immortal Accounts [19], Roach Motel [20, 26, 39]	Privacy [68]
Interface Interference	Cannot delete data from device	Immortal Accounts [19], Roach Motel [20, 26, 39]	Privacy [68]
	No local subscription cancellation	Immortal Accounts [19], Roach Motel [20, 26, 39]	Financial [68]
	General preselection	Preselection [26, 39]	Autonomy [68]
	Visual preference	False Hierarchy [26, 39]	Autonomy [68]
	Confusing text	Trick Question [20, 26, 39]	Autonomy [68]
Subverting Expectations	Confirmshaming	Toying with Emotion [26, 39]	Autonomy [68]
	Forced action	Forced Action [26, 39]	Autonomy [68]
	Hidden information	Hidden Information [26, 39]	Autonomy [68]
	Hidden feature behavior	Aesthetic Manipulation [26, 39]	Cognitive, Financial, Autonomy [68]
	Nagging - General	Nagging [26, 39]	Cognitive [68]
Subverting Expectations	Popup nag	Nagging [26, 39]	Cognitive [68]
	Feature detours to a different modality	Forced Action [26, 39]	Cognitive [68]
	Unprompted suggestions	Nagging [26, 39]	Cognitive, Autonomy [68]
	Nagging self-promotional content	Nagging [26, 39]	Cognitive, Autonomy [68]

Table 2: Final codebook of dark patterns we used to annotate recordings of interactions with IoT devices and companion apps. We group the dark patterns into ten *context categories*, and map each dark pattern to associated traits, strategies, and harms drawn from prior work. Novel dark patterns are shown in bold. Patterns with parenthetical traits or strategies constitute deceptive or unfair behaviors that employ similar strategies to the *maximize privacy dark strategy* [19] but applied to financial or engagement contexts.

this study. To achieve this, we used a video annotation software that facilitated observation and coding for compound or multiple instances of dark patterns, and additionally supported timestamps and image coordinates for each label [28, 29]. When encountering multiple dark patterns on the same screen or video frame, we examine the presented interface elements holistically and consider whether the dark patterns appear to be deployed towards a shared immediate purpose. If so, we select relevant designs in a frame and annotate the selection with multiple dark patterns.

All recordings were annotated by the first author. In cases where the first author felt a label was uncertain, the second author was consulted to achieve consensus.

We validated our device annotations by assessing inter-coder reliability between the first two authors. Specifically, the two authors independently annotated the recordings of one device per type and we compute Cohen's κ to assess agreement. Both authors annotated 585 labels, with 75 and 80 positive unique (binary) labels, and 510 and 505 negative unique labels. For total instance counts, the authors note 169 and 161 positive frequency labels, respectively, and 416 and 424 negative labels per-pattern. For the final dataset, we use the first author's labels by default to maintain consistency across all device annotations.

Table 3 presents the κ agreement statistics with respect to unique and total dark patterns. Across all 52 patterns in our study, we note $\kappa=0.56$ and $\kappa=0.42$ for unique and total dark pattern counts respectively, both of which are in the moderate agreement range ($0.41 < \kappa < 0.60$) [54]. Given the comparatively large size of our codebook and high granularity of individual dark pattern cases, we also grouped our inter-rater labels according to the context categories in our codebook (adapted from Gunawan et al. [42]) and the 16 dark pattern types from Di Geronimo et al. [26].² We observe improvements when calculating κ for each grouped categorization: unique-count $\kappa=0.60$ (mildly significant [54]) for our context categories and $\kappa=0.67$ (significant [54]) for the Di Geronimo et al. [26] categories.³ These results demonstrate more agreement between our labelers at the granularity of categories than at the granularity of specific dark patterns.

In the context of our codebook size, manual methods, video length, and corpus-to-validation sample diversity, we consider our agreement consistent with similar studies' measures [26, 42] and therefore sufficient to proceed with as a reasonable approximation of overall agreement to popular taxonomies. However, as human measurement remains a challenging part of dark patterns study, we further discuss limitations of such methods in subsection 5.4.

4 ANALYSIS

We now analyze our dataset of annotations for all devices included in our experiments. We identify 1,255 total unique instances of dark patterns drawing from 52 distinct patterns. We then compare our results to those from prior measurement studies of dark patterns on the web and in apps.

²These patterns are used in Di Geronimo et al. [26] and Gunawan et al. [42] to group dark pattern cases to the popular Gray et al. [39] taxonomy.

³Whether moderate or significant κ measures are interpreted as acceptable depends on discipline. Within HCI, the adoption of inter-rater reliability measures is somewhat rare [70].

4.1 Dark Patterns Across All Devices

4.1.1 Dark Pattern Popularity. We first count unique dark patterns and per-pattern frequency for all patterns in our codebook across all devices, to broadly quantify dark patterns in IoT experiences.

The cumulative distributive function (CDF) of unique and total dark patterns per device in Figure 2 shows disparity between binary, unique presence counts (whether a dark pattern is found in a device interaction or not) and total frequency counts (how many dark patterns appear in a device interaction, including multiple instances of the same pattern). The x-axis denotes how many dark patterns were discovered (unique count in blue, total in orange). The y-axis represents the percentage of the 57 devices in our study that contained that number of dark patterns for either count. We discovered at least three unique dark patterns in all 57 devices.⁴ On average, devices contained 9 unique dark patterns, and all devices contained < 25 unique patterns.⁵

If the devices in our corpus included only one instance of each unique dark pattern, then the two distributions would be identical and overlaid atop each other. However, beyond the 40th percentile the distributions diverge, with the highest number of total dark patterns (90, Table 4c) being more than triple the maximum number of unique dark patterns (25, Table 4a). Thus, many devices not only exhibit dark patterns multiple times, but do so in large numbers. Table 4 highlights the top ten devices in our corpus by highest and lowest counts of unique and total dark patterns. Both lowest-count tables (Table 4b and Table 4d) share eight out of ten devices. However, the highest-count tables (Table 4a and Table 4c) share only six devices, suggesting variance between top-offending devices' propensity to deploy dark patterns multiple times.

Figure 3 presents the percentage of devices with at least one instance of each dark pattern in our codebook, color-coded by context category. Overall, patterns in the *Interface Interference*, *Consent and Permissions*, *Registration*, *Seen in Settings*, and *Leaving* categories were most frequently adopted.

4.1.2 High Total Counts and Potential Design Templating. As shown in Figure 3, two *Interface Interference* patterns appeared most frequently by total count. A closer look reveals that on average, the *visual preference* pattern appears 6 times per device (the highest average value across all dark patterns), with *general preselection* (the second highest average) appearing only around twice per device—this is visualized in Figure 17 in the Appendix, while Figure 18 stratifies the total count averages per device category to demonstrate that *speakers*, *doorbells*, and *cameras* contained the highest frequencies of these two patterns. We hypothesize that the high adoption rate of these patterns may be due to design templates and/or automated design deployment methods, as opposed to unique, conscious decisions by designers. For example, a UX design tool may have checkbox elements set to have preselected defaults, or binary choice buttons that privilege one button over the other even before button text is added. Context-specific versions of these patterns (e.g., preselected consent checkboxes or settings) were also fairly common in our dataset.

⁴This aligns with prior manual studies' findings of dark patterns in 95% of studied apps [26] and 100% of studied web services [42].

⁵For comparison, prior manual studies note upper-bound counts of 19 unique dark patterns in web services [42] and 23 in apps [26], and average unique counts between 7–8 in both studies.

Stratification	Cohen's κ	
	Unique DPs	Total DPs
Per Pattern	$\kappa=0.561$	$\kappa=0.421$
Per Context Category	$\kappa=0.6$	$\kappa=0.407$
Per DiGeronimo [26] Type	$\kappa=0.679$	$\kappa=0.535$

Table 3: Inter-rater reliability measures computed between the first two authors for unique and total dark patterns, organized per pattern, per context category, and by the 16 types from Di Geronimo et al. [26].




Figure 2: CDFs of unique and total dark patterns per device.

Device	Uniq. №	Device	Uniq. №	Device	Total №	Device	Total №
Fire TV	25	Apple TV	3	Ring Doorbell (*21, Wired)	90	Oxylink Oxygen Monitor	3
Echo Show 5	20	Sengled Smart Hub	3	Fire TV	79	Sengled Smart Hub	3
Withings BPM Connect	19	Oxylink Oxygen Monitor	3	Echo Show 5	75	Withings Sleep	4
Ring Doorbell (*21, Wired)	19	Philips Hue Bridge	3	Nest Mini	69	Amcrest Cam	5
Wyze Cam	18	Withings Sleep	3	Nest Hub Max	57	Homepod Mini	5
Nest Mini	16	Amcrest Cam	4	Ring Camera (Indoor)	55	Thermopro TP90	6
Ring Camera (Indoor)	16	Smartlife LED Bulb	4	Ring Camera	48	GE Microwave	6
Ring Chime	16	Homepod Mini	4	Echo Dot (4th Gen)	38	Smartlife LED Bulb	7
Ring Camera	16	Roku TV	5	Home Mini	38	Roku TV	7
Homepod	15	GE Microwave	6	Nest Doorbell	33	Philips Hue Bridge	8

(a) Highest Unique DPs

(b) Lowest Unique DPs

(c) Highest Total DPs

(d) Lowest Total DPs

Table 4: Top ten IoT devices sorted by those with the (a) highest unique, (b) lowest unique, (c) highest total, and (d) lowest total dark pattern instances.




Figure 3: Percentage of devices containing each dark pattern, with each pattern colored according to our context categories.

4.1.3 Impact of Interaction Duration. To explore the disparity between unique and total counts, we consider the amount of time we spent interacting with each device as a form of robustness check on our methodology: would the disparities disappear if we simply spent more time with each device?

In our interactions, we noticed that device experiences could vary greatly in interface “richness”—an informal measure of available interaction avenues within the experience. This includes feature offerings (e.g., platforms for third-party skills or apps, third-party integrations, built-in analytics or reporting, etc.), device capabilities (e.g., whether a lightbulb is able to control light color or brightness

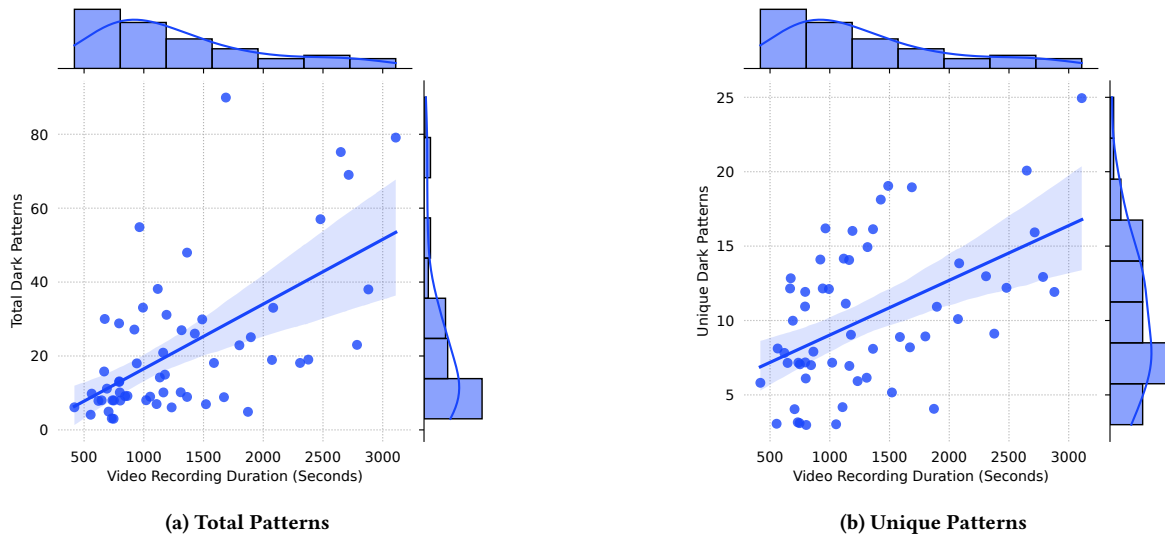


Figure 4: Scatter plots comparing video recording duration to total and unique dark patterns we annotated in each recording. Points are jittered to improve readability. Frequency histograms in both dimensions are shown, as well as a linear regression best-fit with confidence intervals.

instead of just on or off modes), level of detail in options, and design complexity. We posit that interface richness is positively correlated with the number of observed dark patterns: the patterns are design components, so an interface that provides more design surfaces may have greater potential to deploy more dark patterns compared to a leaner interface. Richer interfaces should take longer to traverse experimentally, thus we use device interaction length (represented by video recording length in Table 1) as a proxy measure.

We calculate both Pearson’s and Spearman’s correlation coefficient (r) for recording duration against the number of unique and total dark patterns instances for all devices to find positive and significant correlations⁶ between recording length and dark pattern count: Pearson’s $r = 0.510$ and Spearman’s $r = 0.474$ for unique counts, then Pearson’s $r = 0.591$ and Spearman’s $r = 0.552$ for total counts, with all $p < 0.001$. Figure 4 presents scatter plots of recording duration (x-axis) against total (Figure 4a) or unique (Figure 4b) dark patterns counts (y-axis), with frequency histograms and a linear regression line of best-fit with confidence intervals. For total dark pattern counts, we see tighter clustering around lower counts and outlier behavior (long tails) for high counts, as compared to looser distribution for unique counts. This echoes the two measures’ divergence in Figure 2. However, more research is needed to better understand the relationship between interaction richness and dark pattern deployment, including using models with more robust controls.

4.1.4 Comparison to Prior Measurement Work. Prior modality-specific work measured the presence of unique dark patterns in mobile apps: Di Geronimo et al. [26] investigated 240 apps and

Gunawan et al. [42] inspected 105. Figure 5 compares our findings against these studies, mapped to the context categories from Di Geronimo et al. [26] to provide an apples-to-apples comparison.⁷ Like Gunawan et al. [42], we caution that distributions heavily depend on corpus and codebook construction (which we discuss further in subsection 5.4). Additionally, our interaction methodology departs from both studies’ time-bound interactions, which may impact the discoverability of dark patterns across all these studies.

Our findings generally agree with those from prior work, with some exceptions. Our corpus size is smaller than those in Di Geronimo et al. [26] and Gunawan et al. [42], in part due to IoT companion apps being a strict subset of apps in general, which may explain why we do not observe dark patterns in every Di Geronimo et al. [26] category (e.g., *Hidden Costs*, *Sneaking*, and *Bait & Switch*). We do observe more *Nagging* in our study, potentially because the IoT context presents more opportunities for manufacturers to encourage optional behaviors like linking IoT devices to apps and to each other, or signing-up for optional services. Our study additionally includes more *Hidden Information* and *Trick Questions* patterns compared to Gunawan et al. [42], who only correlate these DiGeronimo categories with one dark pattern each.

4.2 Novel Dark Patterns

In this section we describe newly added dark pattern instances in our codebook (see Table 2) during our pilot experiment and annotation procedures (see section 3), and situate these within extant taxonomies, traits, and strategies.

Some of these novel patterns pertain only to IoT contexts, including: *device sensed without permissions*, *cannot delete data from device*, and *cannot remove device*. Relatedly, we discovered *settings* or

⁶Interpretation ranges for both measures depend on the discipline in question; our measures are considered ‘fair’ in medicine, ‘strong’ in political science, and ‘moderate’ in psychology [8]. We characterize our measures as ‘moderate’ as they fall roughly halfway between no correlation and perfect correlation.

⁷We flatten our findings to one modality in Figure 5 to see how holistic IoT experiences compare to mobile apps, and because the majority of device experiences involved apps or touchscreens.

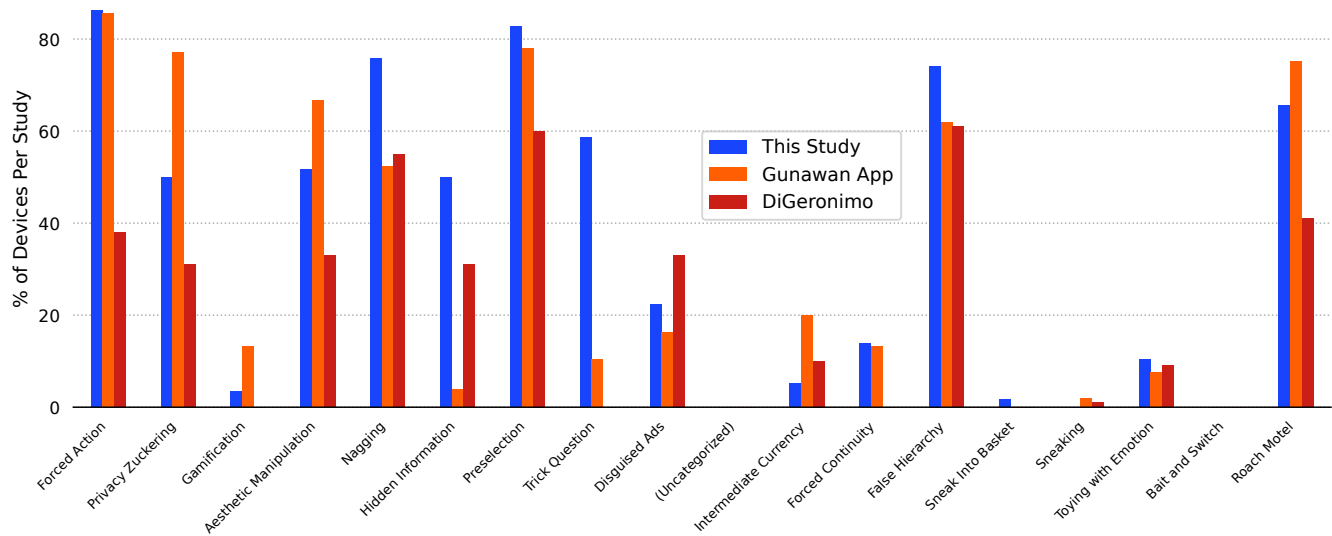


Figure 5: Percentage of devices or app services containing categorized dark patterns, broken down into the Di Geronimo et al. [26] categories and compared against app measurements from prior modality-based studies [42].

features detour to a different modality dark patterns when expected capabilities were not included for a device modality. Prior work generally examined self-contained “experiences” (e.g., websites in a browser or apps on a phone) in isolation rather than as multimodal experiences with configurative dependencies, which may explain why such dark patterns were not previously observed.

Other novel patterns that we identified—such as subscription models or permissions requests—could feasibly be deployed in non-IoT experiences. While permissions- or financial-related dark patterns are not IoT-specific, the additional configuration requirements in IoT devices highlighted novel patterns like *pay for long term use* and *nonpermanent opt-out*.

4.2.1 More Ways to Nag. Nagging [39] patterns manifest in myriad ways, spanning ad-related nags [26] to spammy behavior [20] to extraneous cues [42, 67]. In this study we add two new cases of such patterns, finding *extraneous social media features* and *nagging self-promotional content* in some device experiences.

We considered social media features to be extraneous when they deviated from the primary purpose of a device and were promoted to the user in spammy, aggressive, or otherwise obtrusive manners. Figure 6 presents an example of this pattern that we observed in the Govee Home app for the Govee LED Light Bulb device.

We distinguish *nagging self-promotional content* from native advertising- or shopping-related dark patterns by identifying cases when device experiences or manufacturers presented nags to endorse their own services or content outside of traditional or expected ad placements. These were especially perplexing in the Fire TV as shown in Figure 7: while scrolling for content on a non-Prime user account, we overwhelmingly encountered Prime content carousels with varying promotional labels, and were unable to avoid these carousels. Of the Prime-promoting carousels,

some were labeled as ‘sponsored,’ but the nature of these sponsorships was unclear. Such internal promotions may skirt formal requirements of advertising and disclosure law or guidance. Future research is needed to understand the effect of technically legal but potentially disadvantageous or annoying promotions.

4.2.2 Financial Dark Patterns. The ‘IoT’ part of a smart device is intended to offer consumers value beyond the analog limitations of the device. This presents additional opportunities for manufacturers to apply the financial models from web modalities, like long-term financial relationships via subscription models or tiered access to features. We relate our new case *pay for long term use* to the *Hidden Subscription* [67], *Bait and Switch* [20], and *Obfuscation* dark patterns categories, which we added after being alerted by the Amazon Ring app that certain features were inaccessible due to expired subscriptions during our pilot experiment. Similarly, during our pilot interactions with the Govee app, some features were labeled as exclusive for “Savvy” membership users, but were otherwise accessible to us without signing up for membership, leading to our inclusion of the *feature seems premium but is not* pattern as a counter-case to the previously identified *feature seems free but is not*. We consider these patterns deceptive when they obfuscate key information about device limitations out-of-the-box or upon initial setup, as not all IoT device users may have made the original device purchase.

4.2.3 Settings Inconsistencies. We noted that navigating settings on some devices was particularly challenging or confusing, marking such designs as *inconsistent settings user interfaces*. This falls under *Aesthetic Manipulation*, with particular regards to the ways in which these dark patterns obscure [19], restrict [68], or otherwise interfere with user access to important controls. Figure 8 provides an example from the Aqara Home app, which concurrently shows *preselected notification settings* and *no bulk toggle*. Such designs force users to

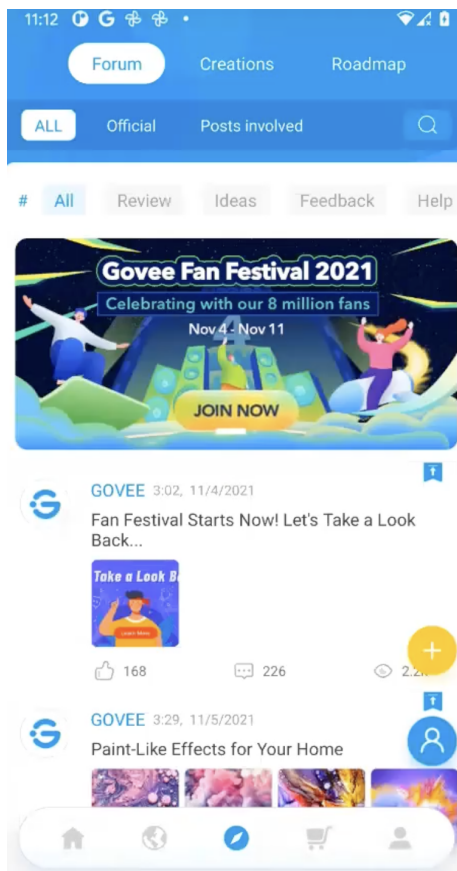



Figure 6: Screenshot of the *extraneous social media features* pattern in the Govee Home app, for the Govee LED Light Bulb device. Prior to this screenshot, the globe menu icon for this page was presented with an *extraneous notification badge* to call user attention to this set of features. The gamut of social media features include user profiles for posting content and interacting with other users, forums, quizzes, gamification, and more.

put in unnecessary additional labor, time, and effort into managing relevant controls, while inconsistencies in the interface potentially prevent users from accessing or understanding available settings.

4.2.4 Detours to Different Modalities. While attempting to use certain settings or features in a device experience, we were sometimes instructed to open an app or visit a website to access that feature. We call these *detours* to describe the manner in which the interface redirected our interactions away from the currently-in-use modality and towards another, with an example of a *settings detour* in Figure 9. Not all features may be feasibly delivered across all modalities in a given service, so we mark this pattern when we are given no notice (prior to attempting to access the setting or feature) that the content is locally unavailable when it should otherwise be reasonable to expect it in the current modality. This dark pattern relates to feature parity issues that have previously been observed across modalities of a singular web service [42], which not only



(a) A clearly marked header for free content in the Fire TV.



(b) A header for sponsored content in the Fire TV. It is unclear whether this content is free, and what the nature of the sponsorship is.



(c) Two headers for Prime content in the Fire TV.

Figure 7: Cropped photographs demonstrating *nagging self-promotional content* in the Amazon Fire TV while users browse for media in a list of thumbnail carousels, and the inconsistent headers found per-carousel. We observed these (among additional instances of this pattern) within seconds of each other, and we were not able to avoid scrolling through Prime-labeled rows while searching for free content.

increase time and labor burdens for users but risk inequitable experiences. Prior work did not assign a label to this behavior; we thus relate *detour* patterns to the *restrict* and *obscure* dark pattern strategies [68].

4.2.5 Device Roach Motels. Prior work [26, 42] articulated several *Roach Motel* [20] cases. We separately define three new device-specific cases (*device sensed without permissions*, *cannot remove device*, and *cannot delete data from device*) to this set of *Roach Motel* patterns in order to investigate the range of possible traps users may encounter in IoT interactions. We noted *devices sensed without permissions* if the app or device itself detected other devices without our active intent to search for other devices within the experiment LAN. One example was the Amazon Alexa app detecting the Amazon Smart Plug device and prompting setup *after* we removed the previously configured device from the app. The latter two patterns are device-specific versions of the *no account deletion*

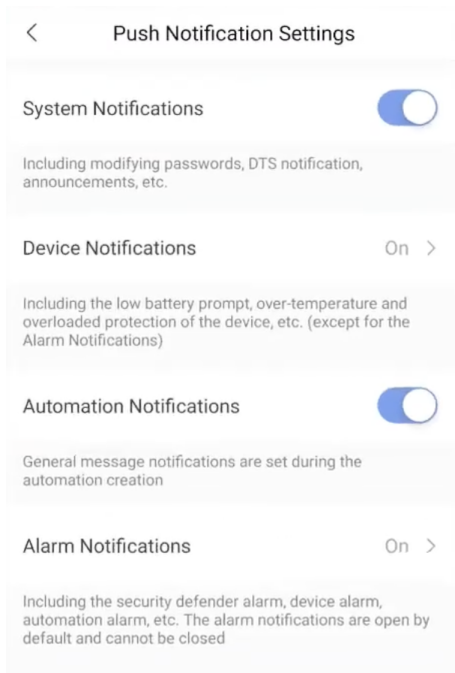


Figure 8: A screenshot of the push notification settings within the Aqara Home app for the Aqara Hub device. These preselected, default-on settings involve both toggles and submenus that have binary on/off indicators, which hinders managing settings in bulk.




Figure 9: A photograph of the settings detour to a different modality pattern from the Fire TV interface. We found this especially interesting as we were prompted to register for Prime earlier in our device interactions.

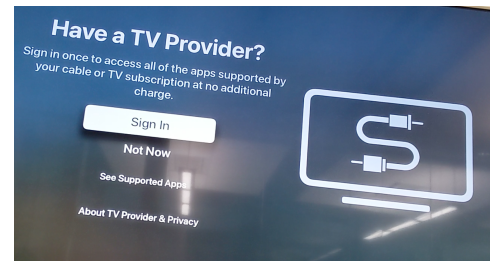
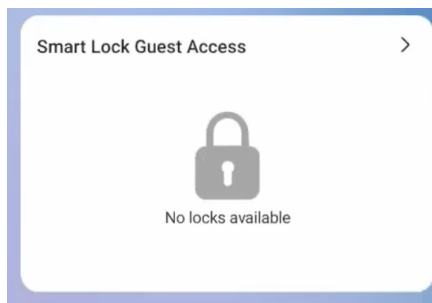
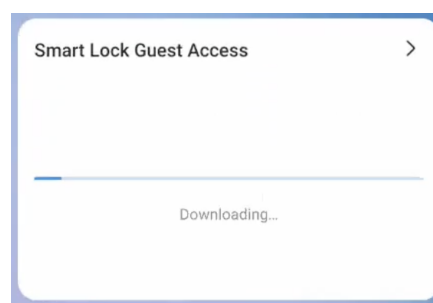


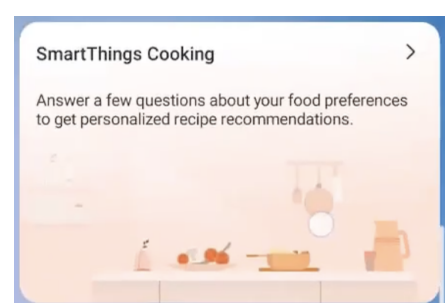
Figure 10: A photograph of the nonpermanent opt-out pattern from the AppleTV interface that appears during setup.



(a) The smart lock feature tile in the SmartThings app, prior to user interaction.



(b) The smart lock feature tile in the SmartThings app, after interaction.



(c) The cooking feature tile in the SmartThings app.

Figure 11: Screenshots demonstrating hidden feature behavior in the Samsung SmartThings companion app, and inconsistencies across similarly-presented designs. The tile in (a) shows no indication that the feature must be downloaded, with (b) showing the download behavior that is triggered by pressing on the tile. The tile in (c) is found below the smart lock feature shown in (a) and (b). Pressing tile (c) triggers a download, then leads to an embedded browser that prompts Samsung account login and integration to a third-party service. The tile’s text offers no indication that additional integrations are required.

or no logout patterns identified in prior work on websites and apps. These patterns help capture different applications of Roach Motels within a single device experience, each with varying degrees of potential impact on user privacy or autonomy. For example, the inability to remove a device from a user account or to delete user data from a device makes it difficult to safely re-sell or gift used

devices. Similarly, devices that are detected by apps prior to users providing specific permissions to the app raise privacy concerns.

4.2.6 Opt-Out Traps. We liken nonpermanent opt-out behavior (shown in Figure 10) to Trick Questions in that such designs can manipulate user interaction towards unintended answers [20]. This dark pattern commonly provides users with binary choices for an

option, with one being an affirmative ‘yes’ and the other being some variation of ‘not now,’ ‘later,’ or ‘skip.’ Users are not provided options like ‘no’ or ‘never.’ Such designs are transparent that the question will be asked again in the future, but do not give users a way to effectively avoid this query or notification and thus accentuates nagging behavior. We label such cases when an interface fails to provide permanent negative options, and do not label cases when users are provided some sort of ‘do not ask again’ choice.

4.2.7 Hidden Feature Behavior. Figure 11 provides an example of *hidden feature behavior*, a sub-case of *hidden information*. We separate this case from the parent dark pattern category for its implications in the IoT context: designs that initially imply that features are inherently provided but later reveal additional conditions for use are thus labeled under this dark pattern. This behavior obfuscates the true value of an IoT service and may deceive consumers into perceiving a product as more robust (sans additional caveats) than it truly is. This pattern is similar in method to feature-related *Money* patterns, which hide the true cost of a feature from the user.

4.2.8 No Local Unsubscribe Options. We added this case following our pilot experiment to investigate whether services with subscription models offered users the ability to cancel subscriptions within the device experience (i.e., without having to head to a website in a browser or otherwise contact the manufacturer). The instances in which this pattern came up were complex: we initially annotated the Ring app as having this pattern as pertaining to the Ring Doorbells and Cameras but not the Ring Chime due to helper text provided to us during attempts to delete devices from the Ring App. Interfaces originally prompting these annotations are shown in Figure 12a and Figure 12b for the Chime and Indoor Camera, respectively. In Figure 12b, the warning message indicates that the subscription-based Ring Protect Plan must be canceled at `ring.com`, thus prompting a *no local unsubscription option* dark pattern. Conversely, the Chime does not include any reference to the Protect Plan and thus we did not mark this dark pattern. However, when we later attempted to delete the entire Ring account via the app, we were presented with the warning message in Figure 12, which explains that any subscriptions on the account will be canceled with account deletion. Due to this technically local unsubscription option, we removed our previous annotations and found no other instances in other devices (seeing as most devices did not mandate or offer subscription models), but retain the dark pattern in our codebook and provide these examples to demonstrate the potential for this behavior in other apps or web services as subscription models are not unique to IoT.

4.3 Dark Patterns by App and Device Modalities

Unlike websites and apps, IoT device experiences span different modalities. Some devices rely on a companion app; others provide direct avenues for interaction via buttons, voice interfaces, and embedded touchscreens. Table 1 presents the companion app dependencies, if any, of the devices in our sample. 72% of the devices in our sample required the use of a companion app, which was surprising for devices like *speakers* (especially the Nest Hub Max, which has an embedded screen) that provide robust interaction methods, and allow many settings to be adjusted sans app.

To examine how dark pattern adoption varies across app-driven and device-driven modalities, we stratify our sample into (a) “App-Only” devices that must be controlled with a companion app at all times ($N = 36$), (b) “Both, For Setup” devices that required an app but only during device setup ($N = 7$), (c) “Both, Non-Setup” devices that required us to use an app after setup to access relevant functionality ($N = 3$), and (d) “Devices Only” devices for which an app is not required and we were never mandated to use an app ($N = 12$). All six devices that required an app for setup were *speakers*. The four devices that prompted app use for non-setup reasons were the Thermopro TP90, the GE Microwave, Renpho Smart Scale, and the Amazon Echo Show 5, and were operational for their primary function prior to our companion app use.

Figure 13 presents the percentage of devices per-modality that included dark patterns, broken down by our dark pattern context categories. Over 60% of devices within all four modalities contained *Consent and Permissions*, *Seen in Settings*, and *Subverting Expectations* dark patterns, making them the most common categories of patterns overall. *Leaving* and *Interface Interference* dark patterns also appeared in over 60% of devices in three modalities and half of the four “Both, Non-setup” devices. The relative popularity of dark patterns in these contexts, even across modalities, speaks to their universal applicability. For example, since these devices are all internet-enabled, consent dialogs are very common, as are settings dialogs that include privacy-sensitive choices.

Dark patterns in the *Registration* category were widely adopted by companion apps, which were used in all categories except “Devices Only”. In particular, we observed that the *account required for use* pattern, which was found to be widely adopted by websites and apps in prior work [42], was widely adopted by IoT devices that required at least some use of a companion app. However, *speakers* (overlapping with the “Both, For Setup” category) revealed an ecosystem-related quirk: because we used a Google Pixel to manage the Google speakers and an iPhone to control HomePod devices, we were already “logged in” to their respective apps and thus were able to use these devices without needing to create another account.⁸ In contrast, devices that did not require or prompt app use (mostly *TVs*, *home appliances*, and *media devices*) treated account creation as optional, i.e., we could use these devices to at least some extent before being asked to create an account.

More rarely adopted categories include *Engagement*, *Ads*, *Money*, and *Shopping*. Of these, adoption of patterns from the *Engagement* category was somewhat consistent, present in around 40% of devices in three modalities. Interpreting the *Ads* category requires understanding our codebook: we generally annotated promotions for third-party content as advertisements, while promotions for first-party, device manufacturer content were considered *self-promotional nags* in the *Subverting Expectations* category. We encountered several *self-promotional nags* across all modalities, but we only observed *native banner ad* designs in *TVs* and *media devices* e.g., the FireTV and LGTV. Dark patterns in the *Money* category were also more prevalent in *media devices* due to requests for free streaming subscriptions and dark patterns that made distinguishing free versus premium content difficult. Adoption of patterns from

⁸Due to this quirk, we did not mark these as requiring account creation specifically for the IoT context. Seamless login behavior like this warrants future scholarship on dark designs within ecosystems or platforms, particularly when single login is used across a manufacturer’s other applications.

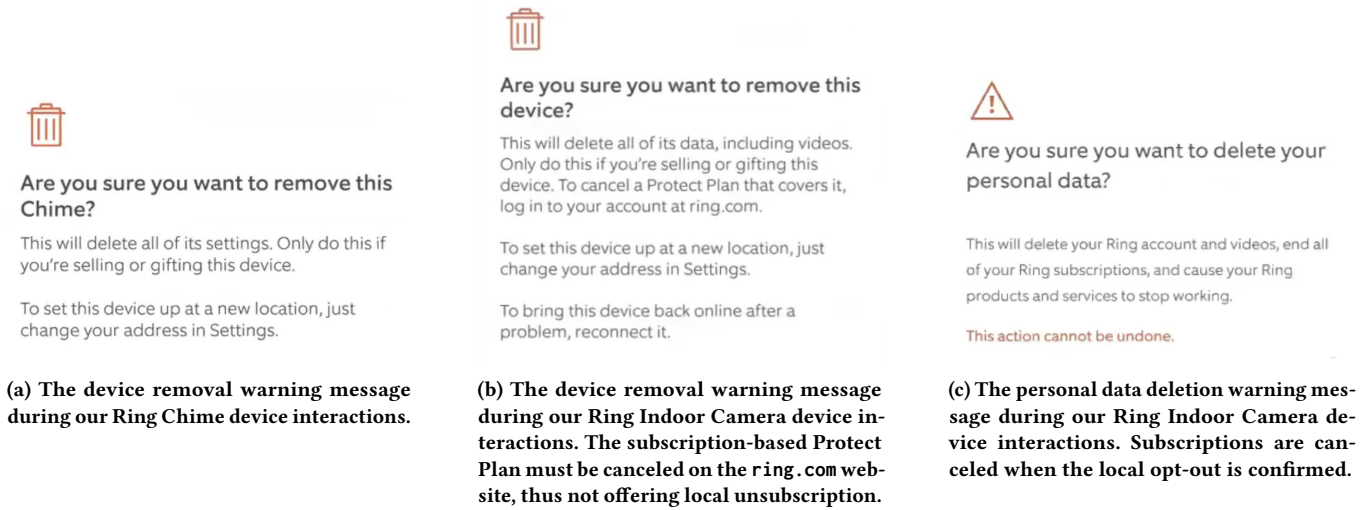


Figure 12: Screenshots from the Ring App, from separate interactions for the (a) Ring Chime and (b,c) Ring Indoor Camera devices. Images (a) and (b) were taken from attempts to remove the device from the Ring App, while (c) was from our attempt to delete the entire Ring account.

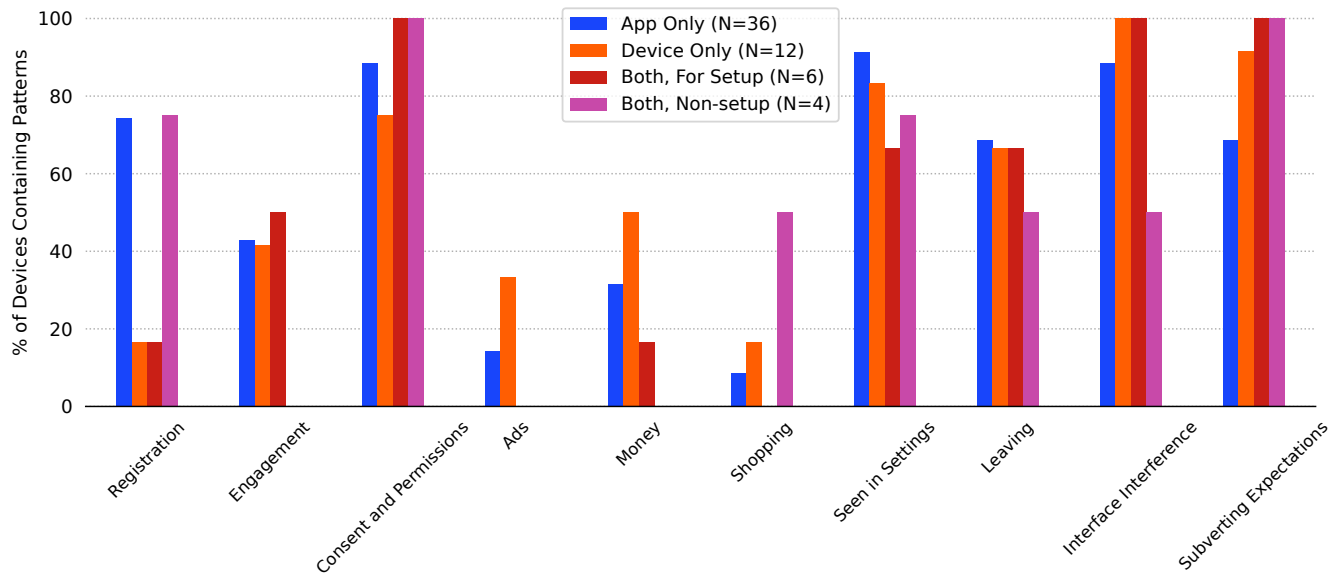


Figure 13: Bar chart showing the percentage of devices containing dark patterns per interaction modality, grouped by dark pattern context category. “App Only” devices could only be controlled by companion apps, while we never used apps to interact with “Device Only” devices. We interacted with ten devices via their physical interface and an app—six out of the seven *speakers* required the companion app only for device setup, while one *speaker* and three other devices directed us to the app later to access other features.

the *Shopping* category was generally low, and we caution against over-interpreting the “Both, For Setup” category in this case due to its small sample size. However, our *Shopping* results align with prior work that also focused on non-shopping centric services [26, 42].

Some IoT experiences are multimodal: device interactions may encompass device hardware, visual displays, and sometimes voice. Speakers, in particular the Amazon Echo Show and Google Nest Hub Max, offer all three modalities as part of one experience: app,

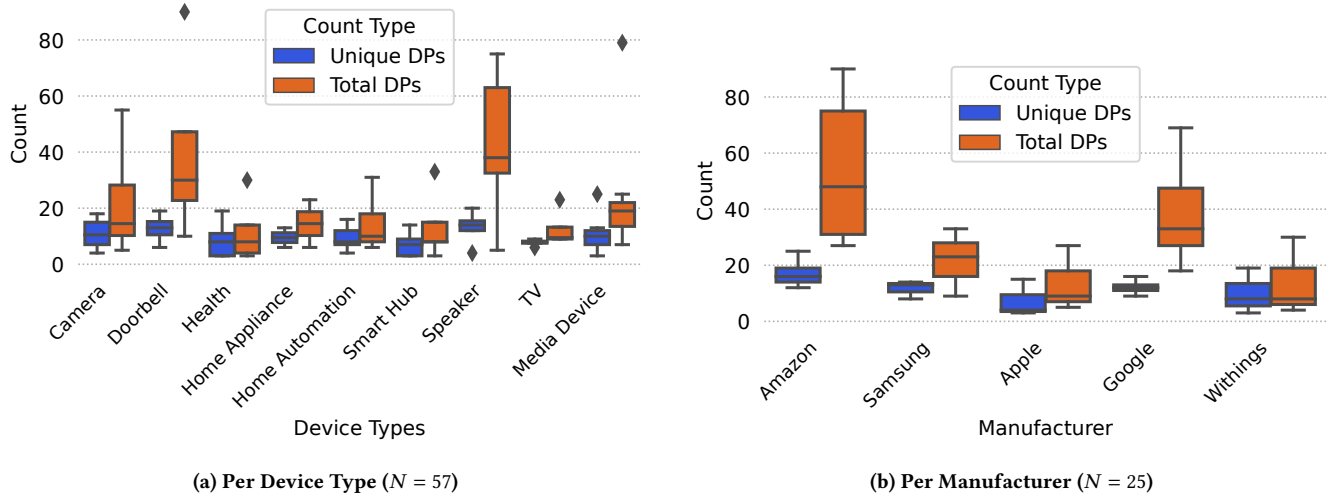


Figure 14: Box plots of unique and total dark pattern instances.

on-device touchscreen, and voice. Every speaker in our test set except the Echo Show 5 required the companion app for setup actions; this was expected for the five speakers without visual interfaces, but surprising for the Google Nest Hub Max, which still mandated the user to install and use respective companion apps despite having a touchscreen that is larger than our Google Pixel smartphone. This constitutes a *forced action* just to use the device at all, and is peculiar for devices that are technologically capable of allowing users to type in registration information on the physical device itself. This behavior also exemplifies our *detour* dark patterns by requiring users to operate multiple modalities for specific needs when the hardware should otherwise afford certain actions. Unlike the Nest Hub Max, the Echo Show 5 (an otherwise functionally similar device), was placed in the “Both, Non-Setup” category. We were able to begin setup without the companion app, but during setup tasks, Alexa automatically ‘checked our app for login’ with alternative options in the device interface. After some time, a skip button unexpectedly appeared. As the device began sensing for the Alexa app without our input and did not provide opt-outs in a timely or transparent fashion, this behavior is reminiscent of *hidden information*, *forced action*, and *bad default* patterns, thus we highlight the differences between these two devices as a demonstration of potentially dark multimodal behavior that researchers might investigate in future work.

4.4 Dark Patterns by Device Type

We now investigate whether dark pattern adoption varies in relation to IoT device type.

Figure 14a presents the distributions of unique and total dark patterns per device type. Three types—*cameras*, *doorbells*, and *speakers*—contained significantly more total dark patterns than other types. The top outlier among *cameras* was the Wansview Cam, containing 61 total dark patterns, while the top outlier among *doorbells* was the Ring Doorbell (*21, Wired), containing 90 total dark patterns (the highest count of any device in our sample). *Speakers*

(representing half of the top ten highest total count devices in Table 4c) yielded the most total dark patterns across device types, with a median of 38 dark patterns. This may be due, in part, to speakers offering at least two interaction modalities (voice and companion app; three in the case of speakers with touchscreens) and requiring longer interactions on our part (four of the top six longest video recordings are *speakers*).

Among the remaining device types, the majority of devices had ≤ 20 unique or total dark patterns. One outlier among *media devices* is Amazon’s Fire TV, which had the highest number of unique dark patterns in our study (25) and 79 total dark pattern instances. Relatedly, while some device types contain more heterogeneous devices than others (e.g., *home appliances* type includes a fridge and a microwave), the Fire TV is one of five relatively fungible streaming devices in the *media devices* type, yet the other four streaming devices’ dark pattern counts do not come close to the Fire TV’s. This disparity is potentially manufacturer-related (see subsection 4.5). Then, compared to the four types with higher counts or extreme outliers, the *health*, *home appliance*, *home automation*, *smart hub*, and *television* types notably contain fewer Amazon, Google, or Apple devices (these manufacturers are not represented at all in *home appliances*, *smart hubs*, or *TVs*).

Figure 15 presents the percentage of devices per type with at least one dark pattern, grouped by our context categories (Table 2). Five categories are thoroughly adopted: there are at least 40% of devices in each device type that contained *Consent and Permission*, *Seen in Settings*, *Leaving*, *Interface Interference*, and *Subverting Expectations* dark patterns. *Registration* and *Engagement* patterns are the runner-ups, but they are adopted less frequently overall than the top five categories and they are not adopted at all by *TVs*. The *TVs* in our sample did not require account registration. *Ads*, *Money*, and *Shopping* patterns were adopted relatively infrequently and inconsistently across device types. This reveals heterogeneity across the IoT ecosystem: perhaps encouragingly, some manufacturers did not insert monetizable content or services into their devices.




Figure 15: Percentage of devices per type that contain ≥ 1 instance of a dark pattern, grouped by dark pattern context category.

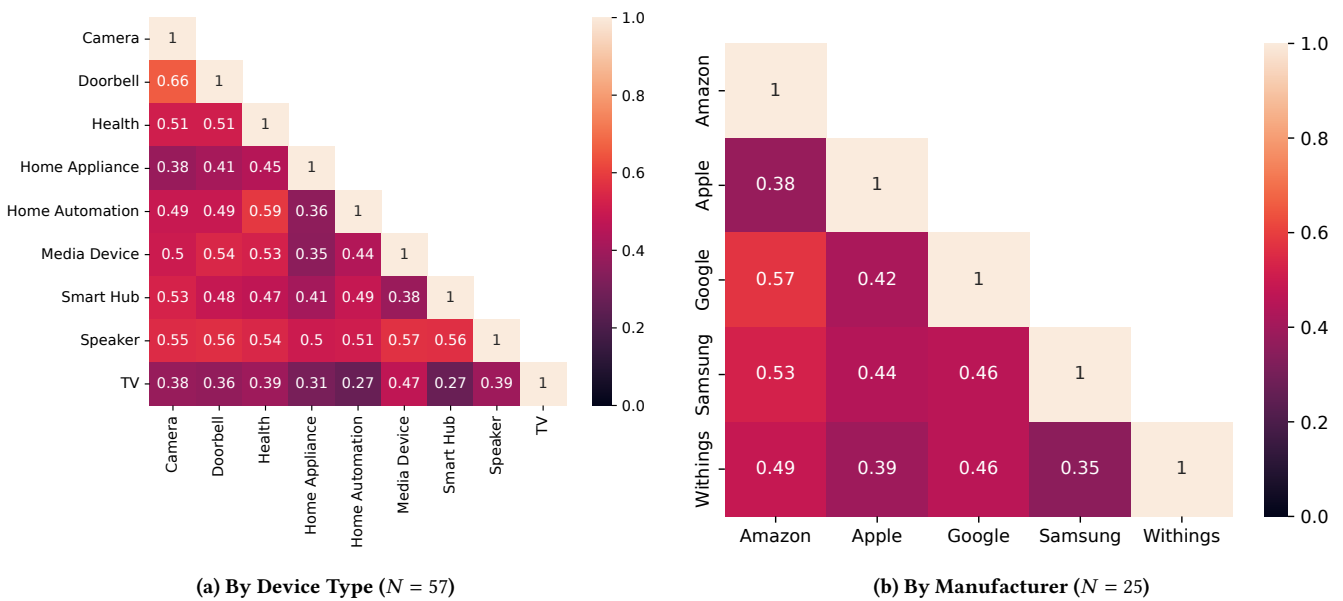


Figure 16: Set overlap of dark patterns. 1 indicates perfect overlap, 0 indicates no overlap.

Though Figure 14a shows that device types were relatively similar in unique dark pattern count across types compared to frequency count, this does not reveal whether the specific unique patterns present among each group of devices overlap. To investigate this, we construct the set of dark patterns found across all devices within each device type, then calculate the Jaccard similarity index between each pair of types, defined as $J(TypeX, TypeY) = |s_X \cap s_Y| / |s_X \cup s_Y|$.

Figure 16a provides the resultant heatmap of these indices, showing generally moderate to low overlap.

The three device-type pairs with the highest overlap are *camera–doorbell* ($J = 0.62$), *home automation–health* ($J = 0.6$), and *speaker–media device* ($J = 0.57$). The relatively high *camera–doorbell* overlap may be explained by the fact that both device types offer similar functionality: all of the doorbells in our study included a camera,

and all interactions with these devices in our study were exclusively conducted through companion apps.

The *home automation* and *health* device categories are both highly diverse: *home automation* devices ranged from light bulbs to thermometers and thermostats, while *health* devices included various monitors to track metrics like sleep quality or blood oxygen. Despite this diversity we find relatively high dark pattern overlap, possibly because devices in both types depended exclusively on companion app interaction. Additionally, such devices are typically engineered for a single purpose (e.g., a blood pressure sleeve only offers features related to blood pressure measurements, while a lightbulb only offers features related to managing light production), offering little opportunity for other features that give rise to unique dark patterns. However, *speakers* exhibited the highest overlap with all other device types (all $J > 0.5$, $\bar{x} = 0.52$). This occurs because *speakers* had the most unique dark patterns overall, thus creating the greatest opportunities for overlap.

Finally, the top five unique dark patterns across all devices per type (listed in Table 7 in the Appendix) tended to be similar across device types, which may account for the lowest overlap in Figure 16a stopping at $J = 0.27$. In particular, *visual preference*, *preselection*, *account requirements*, and *consent- or permissions-related* patterns appear most frequently across device types.

4.5 Dark Patterns by Manufacturer

The IoT market involves a few key manufacturers that produce popular smart devices in contained ecosystems. These include Amazon, Google, Samsung, and Apple, who produce a variety of device types, as well as companies like Withings, who are better known for narrow categories of consumer electronics. 25 devices in our corpus were produced by these five manufacturers, each of whom have three or more devices in our sample (no other manufacturer produced > 2 devices in our corpus and 26 produced one).

Figure 14b presents the distributions of unique and total dark patterns per manufacturer. The median device in our sample contains 15 total dark patterns: Amazon and Google, who produce many *cameras*, *doorbells*, *speakers*, and *media device* devices fall far above the study median (total count medians of 58 and 33 dark patterns, respectively), while other manufacturers fall below it. Figure 16b presents a heatmap of Jaccard indices computed over the sets of dark patterns adopted by pairs of manufacturers and indicates less than 50% overlap between almost all manufacturer pairs, with the exception of *Amazon–Google* and *Amazon–Samsung*, i.e., the three most prolific dark pattern adopters in our sample.

Overall, we do not find clear correlations between company size and dark pattern adoption. Amazon, Apple, Google, and Samsung are some of the largest corporations on earth, yet they do not appear to adopt dark patterns into their products at similar rates. While Amazon products appear the “darkest” in our sample, followed by Google, we caution that our corpus contains fewer samples of devices and different device types from the other three manufacturers. Furthermore, these five manufacturers produce a wide range of device types in our study, which makes it somewhat difficult to conduct apples-to-apples comparisons of dark pattern adoption and frequency for this work.

5 DISCUSSION

In this study, we investigated dark patterns in 57 IoT devices across nine device types by interacting with each device via controlled, scripted experiments. We now discuss our findings and explore our study’s implications for future work and potential dark pattern mitigations.

5.1 How Do IoT Modalities Change Our Understanding of Dark Patterns?

The diversity in IoT devices and their vast applicability to increasing areas of daily life provide an interesting framing for dark patterns measurement. This work demonstrates the multi-factor nature of dark pattern prevalence, and illuminates complexities in the effort to deliver robust IoT services beyond device hardware while managing user expectations for consumer electronics.

5.1.1 Multi-factor Considerations for Dark Pattern Adoption. In this work we examine several potential factors that could influence dark pattern adoption in IoT experiences: interaction context, modality, device type, and manufacturer. We found that visual, screen-based interfaces (e.g., on-device touchscreens and TV screens) drove dark pattern prevalence, rather than companion apps exclusively, non-screen physical device interfaces (e.g., a blood pressure sleeve or lightbulb), or voice modalities (see Figure 13 and Figure 14a).⁹ For example, the Amazon Fire TV contained the largest number of unique dark patterns in our study, but is interacted with only via remote control. Similarly, the fridge, televisions, other media devices, and some smart speakers were fully usable without apps and contained many dark patterns from our codebook. Furthermore, with respect to manufacturers, Amazon and Google adopted more dark patterns among devices in our sample, even when compared to other large manufacturers like Samsung and Apple (see Figure 14b). Overall, however, we found no singular factor that conclusively predicted the presence or lack thereof of dark patterns in IoT devices.

5.1.2 Design (In)Consistencies and Subverted Expectations. In subsection 4.1.2 we note how high dark pattern frequencies may suggest the use of design templates, which can help make UX deployments more efficient at-scale and keep designs consistent (and thus more usable) across a service. This repetition leads users to expect similar behaviors whenever the same design appears. However, this expectation can be subverted to users’ disadvantage. We noted a peculiar example in the SmartThings app; visually consistent designs were used in wildly different manners that made it difficult to determine what features were actually available to us or not out-of-the-box, and to what extent (Figure 11). We considered this to be *hidden feature behavior* as the initial designs offered little transparency into additional requirements or “strings attached.” Such patterns risk deceiving users, particularly for integrations-oriented devices like smart hubs, as users might not fully understand the limitations of a device’s offering at the time of device purchase. Future research on *hidden information*-related patterns could further explore how devices or services communicate or promise value to

⁹In this study, we found voice modalities to be only a minor contributor to dark patterns in our devices when compared to visual modalities, but expect that this is partially due to speaker devices representing a limited portion of our study in device number and methodological focus. Future work should explore voice-controlled interfaces more intently towards uncovering voice dark patterns and prevalence.

end users in their marketing or sales material, towards a goal of improving consumer protections and disclosures.

5.1.3 Add-Ons for Limited Interfaces as Opportunities for Dark Pattern Adoption. Some device types (see subsection 4.4) in our corpus serve single or limited sets of functions but nudge users towards integrations with third-party services. Integrations are also available in app or web services, but one explanation for the increased prevalence of integrations in IoT devices is that doing so can increase the value or functionality of an otherwise purpose-limited device, e.g., by sharing data gathered from device sensors that are not available in other modalities like smartphones. The SmartThings example described in subsection 5.1.2 notes how *hidden feature behavior* might obfuscate real device value to cost users time and money. However, multiple patterns may be involved in nudging users; the Withings app delivered *nagging*, *advertising*, and *hidden information*-related patterns to promote third-party integrations and dedicated an entire page in their app to external health-related subscription services as “Programs” for users to join. We also observe similar behavior from devices produced by large manufacturers (e.g., promoting the manufacturer’s internal services or third-party integrations). Such pattern types were similarly prevalent in prior app studies as demonstrated in Figure 5, especially *nagging* patterns, which were found in over half of apps in both Di Geronimo et al. [26] and Gunawan et al. [42].

Collaborative partnerships can add value to a service. However, failing to disclose whether a feature is included with the price of the device up front may be deceptive. Third-party integrations also raise privacy concerns, especially when the sensitivity of data collected by IoT device sensors is higher than the data accessible to apps and websites.

5.2 Harms Implications from the IoT Context

In addition to harms previously identified in prior work, this study contributes to an understanding of the relationship between device dependencies and darkness. In particular, IoT dark patterns have unique implications for users’ finances, autonomy, and privacy.

5.2.1 Perceived Device Value and Financial Harms. Compared to web- or app-only services, IoT devices necessitate purchase or other means of device ownership, and in some cases require subscriptions in order to ensure continued functionality as described in subsection 4.2. IoT experiences may also attempt to deliver additional services or features to consumers through third-party integrations, some of which may require payment. This complicates user decision making at time of device purchase by obfuscating the true out-of-the-box value of the device, raising financial and autonomy harms for consumers who make purchases with incomplete or incorrect understandings.

5.2.2 Flawed Privacy Controls. IoT devices can change ownership through re-selling or secondhand gifting, exacerbating the IoT privacy and security issues noted in section 2. Dark patterns further add to this problem by interfering with user behavior for important privacy controls. For example, around 60% of devices did not offer account deletion within the visited modalities as demonstrated by Figure 3. With other *Leaving* patterns (e.g., the inability to delete device data or remove a device from an account) as well as the

device reset issues we encountered in subsection 3.1, these dark patterns constitute privacy and autonomy harms by denying users the ability to effectively manage their privacy while using the device and after relinquishing it.

5.2.3 More Roaches in this Motel? Nags for user interaction are not exclusive to the IoT modality, but the increased consumer risks highlighted above may be worsened by nagging behavior. We found such patterns in many of our devices as demonstrated by Figure 3, particularly *non-permanent opt-out* (> 60% of devices), *extraneous notification badges* (> 40% of devices), *nagging self-promotion* (> 40% of devices), with other nagging patterns discovered across multiple context categories. Amazon and Google devices often containing multiple instances of these patterns as noted in the Appendix, Table 8. In light of the vast data collection capabilities of both large manufacturers and their devices, high rates of nagging patterns may influence user behavior towards trackable engagement that leads to further financial or privacy harms. Emergent scholarship is beginning to explore the effect of nags or similar attentional dark patterns on user outcomes [72, 77], though more research is needed to measure dark patterns in attention ecosystems and their impact on other consumer harms.

5.3 Potential Mitigations

Now we explore potential mitigations for the dark patterns findings identified in this work.

5.3.1 Minimize Darkness while Maximizing Value with Design “Appropriateness.” Our work highlights the need for mitigations that minimize dark patterns relative to feature richness, as more complex devices should not be synonymous with more dark patterns. Design practitioners might review design templates and their use to identify possible risks of abusive, deceptive, or unfair applications of those templates. This may help reduce the frequency of designs like those in Figure 11 or popular dark patterns like *non-permanent opt-out* and other *interface interference* patterns including *visual preference* and general *preselection*, which were not only the two most prevalent patterns in our study overall (see Figure 3) but were frequently in the top dark patterns per factor examined in this study (see subsection A.2 for more detail).

One regulatory mitigation of IoT dark patterns could be the idea of “design loyalty” rules [48], which borrow from the law of fiduciary responsibility to prohibit companies from designing their devices, interfaces, and services in a way that conflicts with the best interests of people who use IoT devices. Such rules have already been imposed in California for designs impacting children [90] and proposed in the bipartisan American Data Privacy and Protection Act [2]. One of the significant benefits of loyalty rules is that they can be enforced without requiring strict evidentiary proof of emotional, repetitional, or financial harm. Such harm requirements are a poor fit for dark patterns because of the dispersed, incremental, and often immaterial nature of autonomy and attention-related dangers. Another advantage of loyalty rules is that they direct enforcement agencies to evaluate the relative benefit that flows to companies as well as risks to users. When benefits of a particular design asymmetrically flow toward manufacturers and risks are largely borne by users, the design is disloyal and, thus, dark. Loyalty rules therefore

provide a more structured approach to help lawmakers, companies, and even users identify key choice architectures and understand when design is dark and untrustworthy.

Legislators could additionally borrow “appropriateness” qualities from privacy law’s data minimization principle, which holds that companies should collect, store, and use only data that is adequate, relevant, and limited to what is necessary for a pre-stated purpose [82]. For interactions that are especially important for user controls and autonomy, like consent flows, settings management, or opt-out mechanisms, rules that enforce adequacy and relevancy may improve problems like the use of *detour* patterns or cross-modality equity issues (subsubsection 4.2.4). Design appropriateness perspectives may also address the popularity of *nagging* or engagement-related dark patterns like *extraneous engagement features* (subsubsection 4.2.1), though more research is needed to better understand resulting harms prior to formal regulation of attentional dark patterns.

5.3.2 Build Templates for Desired Design Patterns. Standardization efforts can potentially mitigate dark patterns by providing templates for acceptable design patterns that *promote* autonomy and transparency, countering the effect of templates that lead to consumer harms. One approach is industry standards, inspired by efforts like the Manufacturer Usage Description (MUD) IETF standard [55] that defines the expected behavior of an IoT device and that permits automated compliance tests in deployment [46]. Another approach is to formalize design standards through regulatory rules, similar to the FTC’s .com disclosures [35], the CCPA’s opt-out icon [5], California’s recently enacted Age-Appropriate Design Code Act [90], which provide guidance for avoiding common dark patterns as well as mandates for compliance.

5.3.3 Increase Transparency at Key Interaction Points. Mitigations for dark patterns in the IoT context should also promote improved transparency—particularly at crucial interaction points like account registration and setup—where users are presented with consequential requests (e.g., for account information, permissions, third-party integrations, and other options). Facilitating informed decision-making at the time of device or account configuration can help reduce future privacy or financial risks to the consumer.

Both design and disclosures should offer transparency to consumers, especially so in IoT/consumer electronics experiences as real money and highly sensitive data are at risk. With *Registration* and *Consent and Permissions* dark patterns in > 50% of our devices writ-large (Figure 3) and *Consent and Permissions* issues in > 60% of all devices when stratified by modality (see subsection 4.3 and Figure 13), there is ample room for improvement in transparency surrounding privacy harms. Users must be able to expect that device, feature, or interface functionality corresponds with what they can reasonably infer given the designs or templates they are exposed to. Increasing transparency in feature design can help mitigate consumer harms from issues discussed in subsection 5.1. Lawmakers might consider mandated transparency mechanisms like data protection and design impact assessments, just-in-time disclosure rules, and privacy labeling requirements [4, 35, 82, 90].

5.4 Limitations and Future Work

The methods we used to interact with IoT devices were specifically designed to enable us to explore more design surfaces than an average user might during normal device use, and to do so under carefully controlled, reproducible conditions. Our interaction script does not necessarily reflect the interactions average people would have with IoT devices, and it is not designed to elicit the day-to-day dark patterns they encounter. Under real-world conditions, devices might be moved room-to-room, taken outside the home, linked to other IoT devices, or connected to additional apps to facilitate complex use-cases. Our tests did not include interactions outside of our simulated home environment or between multiple devices and apps concurrently. Further, by creating unique accounts per-device, our study may not reflect the home setup of a person who owns multiple, interconnected IoT devices.

Given these limitations, our measurements should be interpreted as lower-bound estimates of dark pattern prevalence in the IoT devices we tested, and should not be construed as ecologically-valid representations of the types or frequencies of dark patterns that real people might encounter when using IoT devices. Future work could expand our knowledge by taking in-situ measurements of people’s interactions with IoT devices under real-world conditions.

We studied each device during only a single, small window of time. Thus, our interaction approach misses dark patterns that appear only after using a product repeatedly or for extended periods of time. Furthermore, device and companion app behavior may change over time due to firmware and software updates. Future studies may consider taking a longitudinal approach to dark pattern measurements, e.g., to see if experiences get “darker” over time.

Though this study measures dark pattern presence and frequency, it does not discern between higher or lower “darkness” effects between patterns. As such, our results should not be taken as definitive representations of harms severity across devices, but rather as an account of dark design capacity in IoT experiences. The ability to weigh different dark patterns for a more precise measure of dark pattern outcomes is an important area for future research; in particular, scaled measures of dark patterns are necessary for articulating harms towards regulatory thresholds or evidentiary requirements.

Our study is constrained to the devices we had in our possession. While these devices cover many device types and manufacturers, they do not cover them all. Missing device types include: smart watches, rings, and other wearables; cars; and industrial IoT devices. It is unclear whether our results generalize to these other device types or manufacturers.

As discussed in subsubsection 3.4.4 and subsubsection 4.1.4, as well as in prior work [19, 26, 39, 42, 67, 68, 89], manual labeling approaches have weaknesses when attempting to generalize dark patterns research. This can be due to differences in taxonomical interpretation and subjectivity in the concept of ‘darkness.’ User studies similarly demonstrate variance in how participants perceive dark designs. Narrower scopes [19, 26, 67, 89] may yield more clarity but for limited context or for fewer patterns, while broader scopes [39, 68] trade specificity for holistic, dimensional understandings of dark patterns. Our study inspects multifactor, contextual situated IoT devices and thus necessitates a broad approach with a large codebook. Therefore, the generality of our findings is limited

by the same challenges as in other broadly-scoped studies. Future measurement work could delve into specific IoT device types or manufacturers in isolation to trade off breadth for depth, and possibly yield more consistent measurements of dark patterns.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful comments. This research was supported in part by Consumer Reports, a Google ASPIRE award, an NSF grants (#1955227 and #CNS-1900879). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funders.

REFERENCES

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.
- 117th U.S. Congress. 2022. H.R.8152 American Data Privacy and Protection Act. (July 2022). <https://www.congress.gov/bills/117th-congress/house-bill/8152/text>
- Jacob Aagaard, Miria Emma Clausen Knudsen, Per Bækgaard, and Kevin Doherty. 2022. A Game of Dark Patterns: Designing Healthy, Highly-Engaging Mobile Games. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3491101.3519837>
- California Consumer Privacy Act. 2020. California Consumer Privacy Act (Final Text of Proposed Regulations). <https://www.oag.ca.gov/sites/all/files/agweb/pdfs/privacy/oal-sub-final-text-of-regs.pdf>
- California Consumer Privacy Act. 2021. CCPA Opt-Out Icon. <https://oag.ca.gov/privacy/ccpa/icons-download>
- California Privacy Rights Act. 2020. California Privacy Rights Act. https://leginfo.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5
- Imtiaz Ahmad, Rosta Farzan, Apu Kapadia, and Adam J. Lee. 2020. Tangible Privacy: Towards User-Centric Sensor Designs for Bystander Privacy. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2 (oct 2020). <https://doi.org/10.1145/3415187>
- Haldun Akgözü. 2018. User's guide to correlation coefficients. *Turkish Journal of Emergency Medicine* (2018). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6107969/>
- Noura Aleisa and Karen Renaud. 2017. Privacy of the Internet of Things. *Proceedings of the 50th Hawaii International Conference on System Sciences* (Jan 2017), 5947–5956. <https://doi.org/10.24251/HICSS.2017.717>
- Omar Alrawi, Chaz Lever, Manos Antonakakis, and Fabian Monrose. 2019. SoK: Security Evaluation of Home-Based IoT Deployments. In *2019 IEEE Symposium on Security and Privacy (SP)*, 1362–1380. <https://doi.org/10.1109/SP.2019.00013>
- Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, et al. 2017. Understanding the Mirai Botnet. In *Proc. of USENIX Security Symposium*, 1093–1110.
- Noah Athorpe, Dillon Reisman, and Nick Feamster. 2016. A Smart Home is No Castle: Privacy Vulnerabilities of Encrypted IoT Traffic. *Data and Algorithmic Transparency Workshop* (2016), 6.
- Daniel Arp, Erwin Quiring, Christian Wressnegger, and Konrad Rieck. 2017. Privacy Threats through Ultrasonic Side Channels on Mobile Devices. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*, 35–47.
- Natá Micael Barbosa, Joon S Park, Yaxing Yao, and Yang Wang. 2019. "What if?" Predicting Individual Users' Smart Home Privacy Preferences and Their Changes. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019), 211–231.
- Rubem Barbosa-Hughes. 2020. Interaction Patterns using Machine Learning and Location Services in User Interfaces for the Consumer IoT. In *Proceedings of the European Conference on Pattern Languages of Programs 2020 (EuroPLoP '20)*. Association for Computing Machinery, 1–9. <https://doi.org/10.1145/3424771.3424777>
- Luiz Adolpho Baroni, Allison Andrey Puska, Luciana Cardoso de Castro Salgado, and Roberto Pereira. 2021. *Dark Patterns: Towards a Socio-Technical Approach*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3472301.3484336>
- Dieter Bohn. 2017. Google's Home Mini needed a software patch to stop some of them from recording everything. <https://www.theverge.com/2017/10/10/16456050/google-home-mini-always-recording-bug>
- Kerstin Bongard-Blanchy, Ariana Rossi, Salvador Rivas, Sophie Doublet, Vincent Koening, and Gabriele Lenzini. 2021. "I am Definitely Manipulated, Even When I am Aware of it. It's Ridiculous!" - Dark Patterns from the End-User Perspective. In *Designing Interactive Systems Conference 2021*. <http://doi.org/10.1145/3461778.3462086>
- Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns. *Proc. of PETS* 2016, 4 (2016), 237–254. <https://content.sciendo.com/view/journals/popets/2016/4/article-p237.xml>
- Harry Brignull. 2010. Types of Deceptive Design. <https://www.deceptive.design/types>.
- Bringing Dark Patterns to Light: An FTC Workshop 2021. <https://www.ftc.gov/news-events/events-calendar/bringing-dark-patterns-light-ftc-workshop>
- Z Berkay Celik, Leonardo Babun, Amit Kumar Sikder, Hidayet Aksu, Gang Tan, Patrick McDaniel, and A Selcuk Uluagac. 2018. Sensitive Information Tracking in Commodity IoT. In *27th USENIX Security Symposium (USENIX Security 18)*, 1687–1704.
- Camille Cobb, Sruti Bhagavatula, Kalil Anderson Garrett, Alison Hoffman, Varun Rao, and Lujo Bauer. 2021. "I would have to evaluate their objections": Privacy tensions between smart home device owners and incidental users. *Proceedings on Privacy Enhancing Technologies* 2021, 4 (2021), 54–75.
- Competition and Markets Authority of the United Kingdom. 2022. *Evidence Review of On-line Choice Architecture and Consumer and Competition Harm*. Technical Report 157, 261 pages. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1069423/OCA_Evidence_Review_Paper_14_4_22.pdf
- Gregory Conti and Edward Sobieski. 2010. Malicious Interface Design: Exploiting the User. In *Proc. of WWW*.
- Linda Di Geronimo, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. UI Dark Patterns and Where to Find Them: A Study on Mobile Applications and User Perception. In *Proc. of CHI*.
- Daniel J. Dubois, Roman Kolcun, Anna Maria Mandalari, Muhammad Talha Paracha, David Choffnes, and Hamed Haddadi. 2020. When Speakers Are All Ears: Characterizing Misappropriations of IoT Smart Speakers. In *Proc. of the Privacy Enhancing Technologies Symposium (PETS)*.
- A. Dutta, A. Gupta, and A. Zissermann. 2016. VGG Image Annotator (VIA). <http://www.robots.ox.ac.uk/vgg/software/via/>.
- Abhishek Dutta and Andrew Zisserman. 2019. The VIA Annotation Software for Images, Audio and Video. In *Proceedings of the 27th ACM International Conference on Multimedia (Nice, France) (MM '19)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3343031.3350535>
- Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2021. Which Privacy and Security Attributes Most Impact Consumers' Risk Perception and Willingness to Purchase IoT Devices?. In *2021 IEEE Symposium on Security and Privacy (SP)*, IEEE, 519–536. <https://doi.org/10.1109/SP40001.2021.00112>
- Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. 2019. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 1–12. <https://doi.org/10.1145/3290605.3300764>
- Earlene Fernandes, Jaeyoon Jung, and Atul Prakash. 2016. Security analysis of emerging smart home applications. In *2016 IEEE symposium on security and privacy (SP)*. IEEE, 636–654.
- Organization for Economic Co-operation and Development. 2022. Dark Commercial Patterns. <https://www.oecd-ilibrary.org/docserver/44f5e846-en.pdf?expires=1670358058&id=id&accname=guest&checksum=C06C85A96CF08E11AD88986E583DF54>
- Directorate-General for Justice and Consumers (European Commission), Francisco Lupiáñez-Villanueva, Alba Boluda, Francesco Bogliacino, Giovanni Liva, Lucie Lecharroy, and Teresa Rodriguez de las Heras Ballell. 2022. *Behavioural study on unfair commercial practices in the digital environment: dark patterns and manipulative personalisation : final report*. Technical Report. Publications Office of the European Union, LU. <https://data.europa.eu/doi/10.2838/859030>
- FTC. 2013. .com Disclosures. <https://www.ftc.gov/sites/default/files/attachments/press-releases/ftc-staff-revises-online-advertising-disclosure-guidelines/130312dotcomdisclosures.pdf>
- Google. 2022. Explore what you can do with Google Nest or Home devices. <https://support.google.com/googlenest/answer/7130274>
- Colin M. Gray. 2019. The dark side of UX Design. <https://darkpatterns.uxp2.com/>.
- Colin M. Gray, Jingle Chen, Shruthi Sai Chivukula, and Liyang Qu. 2021. End User Accounts of Dark Patterns as Felt Manipulation. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct 2021), 372:1–372:25. <https://doi.org/10.1145/3479516>
- Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. 2018. The Dark (Patterns) Side of UX Design. In *Proc. of CHI*.
- Colin M. Gray, Cristiana Santos, Natalia 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 (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 172, 18 pages. <https://doi.org/10.1145/3411764.3445779>
- Paul Graßl, Hanna Schraffenberger, Frederik Zuiderveen Borgesius, and Moniek Buijzen. 2021. Dark and Bright Patterns in Cookie Consent Requests. *Journal of Digital Social Research* 3, 1 (Feb. 2021), 1–38. <https://doi.org/10.33621/jdsr.v3i1.54> Number: 1.
- Johanna Gunawan, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2021. A Comparative Study of Dark Patterns Across Mobile and Web Modalities. (Oct. 2021).
- Johanna Gunawan, Cristiana Santos, and Irene Kamara. 2022. Redress for Dark Patterns Privacy Harms? A Case Study on Consent Interactions. In *Proceedings of the 2022 Symposium on Computer Science and Law (CSLAW '22)* (Washington, D.C., USA) (CSLAW '22). Association for Computing Machinery.
- Frøde Guribye, Oda Elise Nordberg, Are Nyhammer, Marija Slavkovic, and Than Htut Soe. 2021. Dark patterns in cookie consent notices: new definitions and mitigation strategies. *Position Papers of CHI'22 "What Can CHI Do About Dark Patterns?"* (2021). https://drive.google.com/file/d/1V-P98rH_wqLzLaqY1HZPPGfHcCKWbX/view
- Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. 2022. "Okay, whatever": An Evaluation of Cookie Consent Interfaces. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–27. <https://doi.org/10.1145/3491102.3501985>
- Ayyoob Hamza, Dinesha Ranathunga, Hassan Habibi Gharakheili, Theophilus A. Benson, Matthew Roughtan, and Vijay Sivaraman. 2022. Verifying and Monitoring IoTs Network Behavior Using MUD Profiles. *IEEE Transactions on Dependable and Secure Computing* 19, 1 (2022), 1–18.
- Woodrow Hartzog. 2018. *Privacy's Blueprint: The Battle to Control the Design of New Technologies*. Harvard University Press.
- Woodrow Hartzog and Neil M. Richards. 2022. Legislating Data Loyalty. *Notre Dame Law Review Reflection* 97 (June 2022). Issue 356. <https://ssrn.com/abstract=4131523>
- Umar Iqbal, Pounneh Nikkha Bahrami, Rahmadi Trimananda, Hao Cui, Alexander Gamero-Garrido, Daniel Dubois, David Choffnes, Athina Markopoulou, Franziska Roesner, and Zubair Shafiq. 2022. Your Echos are Heard: Tracking, Profiling, and Ad Targeting in the Amazon Smart Speaker Ecosystem. <https://arxiv.org/abs/2204.10920>
- Dave Johnson and William Antonelli. 2021. The most important Alexa voice commands you can use with your Amazon Echo. <https://www.businessinsider.com/guides/tech/alexa-commands>

- [51] Jennifer King and Adriana Stephan. 2021. Regulating Privacy Dark Patterns in Practice - Drawing Inspiration from the California Privacy Rights Act. *Georgetown Law Technology Review* 5 (2021), 26 pages. Issue 250.
- [52] Chiara Krisam, Heike Dietmann, Melanie Volkamer, and Oksana Kulyk. 2021. Dark Patterns in the Wild: Review of Cookie Disclaimer Designs on Top 500 German Websites. In *European Symposium on Usable Security 2021* (Karlsruhe, Germany) (*EuroUSEC '21*). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3481357.3481516>
- [53] Jacob Leon Kröger, Leon Gellrich, Sebastian Pape, Saba Rebecca Brause, and Stefan Ullrich. 2022. Personal information inference from voice recordings: User awareness and privacy concerns. *Proceedings on Privacy Enhancing Technologies* 2022, 1 (2022), 6–27.
- [54] J. Richard Landis and Gary G. Koch. 1977. The Measurement of Observer Agreement for Categorical Data. Issue 1. <https://doi.org/10.2307/2529310>
- [55] Eliot Lear, Ralph Droms, and Dan Romascanu. 2019. Manufacturer Usage Description Specification. RFC 8520. <https://www.rfc-editor.org/info/rfc8520>
- [56] Neha Lingareddy, Brennan Schaffner, and Marshini Chetty. 2021. Can I Delete My Account? Dark Patterns in Account Deletion on Social Media. *Position Papers of CHI'22 "What Can CHI Do About Dark Patterns?"* (2021). https://www.google.com/url?q=https://drive.google.com/file/d/190uWk2bhJngE0J_K_kGt2v91z_f010cz/view?usp%3Dsharing&sa=D&source=editor&ust=1619579203736000&usq=AFQjCNHhFRvNyZ0NrYr7AVXUqAKvIKd-tQ
- [57] Haoyu Liu, Tom Spink, and Paul Patras. 2019. Uncovering security vulnerabilities in the Belkin WeMo home automation ecosystem. In *2019 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 894–899.
- [58] Natasha Lomas. 2022. Amazon agrees to drop Prime cancellation 'dark patterns' in Europe. <https://techcrunch.com/2022/07/01/amazon-ends-prime-cancellation-dark-patterns-europe/>
- [59] Kai Lukoff, Alexis Hiniker, Colin M. Gray, Arunesh Mathur, and Shruthi Sai Chivukula. 2021. *What Can CHI Do About Dark Patterns?* Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3411763.3411360>
- [60] Aditi M. Bhoat, Mayuri A. Shinde, and Wricha P. Mishra. 2020. Towards the Identification of Dark Patterns: An Analysis Based on End-User Reactions. In *IndiaHCI '20: Proceedings of the 11th Indian Conference on Human-Computer Interaction* (Online, India) (*IndiaHCI 2020*). Association for Computing Machinery, New York, NY, USA, 24–33. <https://doi.org/10.1145/3429290.3429293>
- [61] Eryn Ma and Eleanor Birrell. 2022. Prospective Consent: The Effect of Framing on Cookie Consent Decisions. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3491101.3519687>
- [62] Nathan Malkin, Joe Deatrack, Allen Tong, Primal Wijesekera, Serge Egelman, and David Wagner. 2019. Privacy attitudes of smart speaker users. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019).
- [63] Anna Maria Mandalari, Daniel J. Dubois, Roman Kolcun, Muhammad Talha Paracha, Hamed Haddadi, and David Choffnes. 2021. Blocking without Breaking: Identification and Mitigation of Non-Essential IoT Traffic. In *Proc. of the Privacy Enhancing Technologies Symposium (PETS)*.
- [64] Shirang Mare, Franziska Roesner, and Tadayoshi Kohno. 2020. Smart Devices in Airbnbs: Considering Privacy and Security for both Guests and Hosts. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (2020), 436–458.
- [65] Karola Marky, Nina Gerber, Michelle Gabriela Pelzer, Mohamed Khamis, and Max Mühlhäuser. 2022. "You offer privacy like you offer tea": Investigating Mechanisms for Improving Guest Privacy in IoT-Equipped Households. *Proceedings on Privacy Enhancing Technologies* 1 (2022), 21.
- [66] Karola Marky, Verena Zimmermann, Alina Stöver, Philipp Hoffmann, Kai Kunze, and Max Mühlhäuser. 2020. All in One! User Perceptions on Centralized IoT Privacy Settings (*CHI EA '20*). Association for Computing Machinery, 1–8. <https://doi.org/10.1145/3334480.3333016>
- [67] Arunesh Mathur, Gunes Acar, Michael Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW (2019).
- [68] Arunesh Mathur, Jonathan Mayer, and Kihir Kshirsagar. 2021. What Makes a Dark Pattern... Dark? Design Attributes, Normative Considerations, and Measurement Methods. <https://arxiv.org/pdf/2101.04843.pdf>
- [69] M. Hammad Mazhar and Zubair Shafiq. 2020. Characterizing Smart Home IoT Traffic in the Wild. *ACM/IEEE Conference on Internet of Things Design and Implementation* (Mar 2020). <http://arxiv.org/abs/2001.08288>
- [70] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-Rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72 (nov 2019), 23 pages. <https://doi.org/10.1145/3359174>
- [71] Thomas Mildner and Gian-Luca Savino. 2021. *Ethical User Interfaces: Exploring the Effects of Dark Patterns on Facebook*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3411763.3451659>
- [72] Alberto Monge Roffarello and Luigi De Russis. 2022. Towards Understanding the Dark Patterns That Steal Our Attention. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3491101.3519829>
- [73] Sara Morrison. 2021. Dark patterns, the tricks websites use to make you say yes, explained. <https://www.vox.com/recode/22351108/dark-patterns-ui-web-design-privacy>
- [74] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Proc. of USENIX Security Symposium*. 399–412. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/naeini>
- [75] Arvind Narayanan, Arunesh Mathur, Marshini Chetty, and Mihir Kshirsagar. 2020. Dark Patterns: Past, Present, and Future. *ACM Queue* 18, 2 (2020).
- [76] Commission nationale de l'informatique et des libertés (CNIL). 2022. Cookies: the CNIL fines GOOGLE a total of 150 million euros and FACEBOOK 60 million euros for non-compliance with French legislation. (January 2022). <https://www.cnil.fr/en/cookies-cnil-fines-google-total-150-million-euros-and-facebook-60-million-euros-non-compliance>
- [77] Aileen Nielsen. 2022. Tech has an Attention Problem. *Berkeley Center for Long-Term Cybersecurity Whitepapers* (2022). https://cltc.berkeley.edu/wp-content/uploads/2021/09/CLTC_Techs_Attention_Problem.pdf
- [78] 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 *Proc. of CHI*.
- [79] SIG on Governing the Social Media and Data Economy. 2022. Futureproof Methods for Measuring and Detecting Dark Patterns. <https://www.uu.nl/en/events/future-proof-methods-for-measuring-and-detecting-dark-patterns>
- [80] Kentrell Owens, Johanna Gunawan, David Choffnes, Pardis Emami-Naeini, Tadayoshi Kohno, and Franziska Roesner. 2022. Exploring Deceptive Design Patterns in Voice Interfaces. In *2022 European Symposium on Usable Security* (Karlsruhe, Germany) (*EuroUSEC 2022*). Association for Computing Machinery, New York, NY, USA, 64–78. <https://doi.org/10.1145/3549015.3554213>
- [81] Muhammad Talha Paracha, Daniel J Dubois, Narseo Vallina-Rodriguez, and David Choffnes. 2021. IoTLS: Understanding TLS Usage in Consumer IoT Devices. In *Proc. of the Internet Measurement Conference*.
- [82] European Parliament and Council of European Union. 2016. EU General Data Protection Regulation (GDPR). <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN>
- [83] David Priest, Tauren Dyson, and Taylor Martin. 2022. Every Alexa command you can give your speaker. <https://www.cnet.com/how-to/every-alexa-command-you-can-give-your-amazon-echo-smart-speaker/>
- [84] Eric Ravenscraft. 2020. How to Spot—and Avoid—Dark Patterns on the Web. <https://www.wired.com/story/how-to-spot-avoid-dark-patterns/>
- [85] Jingjing Ren, Daniel J. Dubois, David Choffnes, Anna Maria Mandalari, Roman Kolcun, and Hamed Haddadi. 2019. Information Exposure for Consumer IoT Devices: A Multidimensional, Network-Informed Measurement Approach. In *Proc. of the Internet Measurement Conference (IMC)*.
- [86] Eyal Ronen and Adi Shamir. 2016. Extended functionality attacks on IoT devices: The case of smart lights. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 3–12.
- [87] Eyal Ronen, Adi Shamir, Achi-Or Weingarten, and Colin O'Flynn. 2017. IoT goes nuclear: Creating a ZigBee chain reaction. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 195–212.
- [88] Said Jawad Saidi, Anna Maria Mandalari, Roman Kolcun, Hamed Haddadi, Daniel J. Dubois, David Choffnes, Georgios Smaragdakis, and Anja Feldmann. 2020. A Haystack Full of Needles: Scalable Detection of IoT Devices in the Wild. In *Proc. of the Internet Measurement Conference (IMC)*.
- [89] 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, Article 417 (nov 2022), 43 pages. <https://doi.org/10.1145/3555142>
- [90] California Legislature (2021-2022 Regular Session). 2022. California Age-Appropriate Design Code Act. https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=202120220AB2273
- [91] Stanford Digital Civil Society. 2021. Dark Patterns Tip Line. <https://darkpatternstipline.org>
- [92] Than Htut Soe, Oda Elise Nordberg, Frode Guribye, and Marija Slavkovic. 2020. *Circumvention by Design - Dark Patterns in Cookie Consent for Online News Outlets*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3419249.3420132>
- [93] Daniel Susser, Beate Roessler, and Helen Nissenbaum. 2019. Online Manipulation: Hidden Influences in a Digital World.
- [94] Richard H. Thaler and Cass R. Sunstein. 2008. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press.
- [95] Janus Varmarken, Hieu Le, Anastasia Shuba, Athina Markopoulou, and Zubair Shafiq. 2020. The TV is Smart and Full of Trackers: Measuring Smart TV Advertising and Tracking. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (Apr 2020), 129–154. <https://doi.org/10.2478/popets-2020-0021>
- [96] Megan Wollerton. 2020. Wink's subscription plan kicks off July 27. <https://www.cnet.com/home/smart-home/winks-subscription-plan-kicks-off-july-27/>
- [97] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata McDonough, and Yang Wang. 2019. Privacy Perceptions and Designs of Bystanders in Smart Homes. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (nov 2019). <https://doi.org/10.1145/3359161>
- [98] Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyuan Xu. 2017. Dolphinattack: Inaudible voice commands. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 103–117.
- [99] Serena Zheng, Noah Aporthe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (Nov 2018), 200:1–200:20. <https://doi.org/10.1145/3274469>
- [100] Wei Zhou, Yan Jia, Yao Yao, Lipeng Zhu, Le Guan, Yuhang Mao, Peng Liu, and Yuting Zhang. 2019. Discovering and Understanding the Security Hazards in the Interactions between IoT Devices, Mobile Apps, and Clouds on Smart Home Platforms. In *28th USENIX Security Symposium (USENIX Security 19)*. 1133–1150.
- [101] Mohammed Zubair, Devrim Unal, Abdulla Al-Ali, and Abdullatif Shikfa. 2019. Exploiting Bluetooth Vulnerabilities in e-Health IoT Devices. In *Proceedings of the 3rd International Conference on Future Networks and Distributed Systems (ICFNDS '19)*. Association for Computing Machinery, 1–7. <https://doi.org/10.1145/3341325.3342000>

A APPENDIX

This appendix provides additional figures, tables, and analysis that supplement the main body of work in this paper.

A.1 Supplementary Device Information

Device Type	Device Name	Device Firmware Version	App Name	App Version
Home Automation	Amazon Smart Plug	Up to date (11/5/21)	Amazon Alexa	2.2.432925.0
	Jinwoo Smart Bulb	Up to date (6/10/22)	Jinwoo Smart	2.0.9
	Gosund Smart Light Bulb	V3.3.35	Gosund	4.5.0
	Govee LED Light Bulb	Up to date (11/8/21)	Govee Home	4.5.5
	Magichome Strip	1.8.1	Magic Home Pro	33.v4.17.6445-A
	Meross Door Opener	3.2.5	meross	2.32.4
	Nest Thermostat*	6.2-22	Nest	5.67.0.6
	Ring Chime	Up to date (11/5/2021)	Ring	3.45.1
	Smartlife LED Bulb	3.32.5	Smart Life	2.3.0
	WeMo Plug	Up to date (11/4/21)	Wemo	1.29.1
	Thermopro TP90	Up to date (11/4/21)	ThermoPro Home	1.0.4
	TP-Link Bulb	1.8.11	Kasa Smart	2.35.0.1021
	TP-Link Plug	1.2.6	Kasa Smart	2.35.0.1021
	Camera	Amcrest Cam	V2.400.AC02.15.R	Amcrest View Pro
Arlo Q Cam		Up to date (6/10/22)	Arlo Secure: Home Security	Ver Jun 6 2022
D-Link Cam		2.06.03(3.5.18-b01)	mydlink	2.6.1
Lefun Cam		Up to date (11/01/21)	MIPC	v8.9.3.2109291620
Nest Camera		Up to date (6/13/22)	Google Home	Ver Jun 8 2022
Ring Camera		Up to date (12/13/2021)	Ring	3.46.0
Ring Camera (Indoor)		Up to date (6/10/22)	Ring	Ver Jun 8 2022
Tuya Smart Camera		V5.2.7 (Main Module) V5.2.7 (MCU Module)	Tuya Smart	3.32.5
Wyze Cam		4.36.6.17	Wyze	2.25.31
Yi Home Camera		2.1.0.0E_201809191630	Yi Home	5.1.4_20211020
Media Device	Apple TV*	tvOS 15.0 (19J346)		
	Chromecast w/ Google TV*	4.9.180 (Kernel version) 10 (Android TV OS version) QTS1.210311.008.7350836 (Android TV OS Build)		
	Facebook Portal Mini*	1.28.1 (Software version)		
	Fire TV*	Fire OS 5.2.8.4(672751320) (Software Version) 6330056.1 (Fire TV Home Version)		
	Nintendo Switch*	13.1.0		
	Roku TV*	10.0.0 (Software version) 4209 (Build)		
Smart Hub	TiVo Stream*	9 (OS version) 1.0.902-53 (App version) 1.4.191 (TiVo+ version)		
	Aqara Hub	3.3.2_0010.0610	Aqara Home	2.3.17
	Sengled Smart Hub	Up to date (11/23/21)	Sengled Home	2.1.9
	SmartThings Hub	000.039.00006	SmartThings	1.7.73.22
	Switchbot Hub	Up to date (11/3/21)	SwitchBot	5.4.0.8
Doorbell	Philips Hue Bridge	1.47.1947108030 (Software)	Philips Hue	4.7.0
	Arlo Doorbell	Up to date (6/10/22)	Arlo Secure: Home Security	Ver Jun 6 2022
	Nest Doorbell	Up to date (6/13/22)	Google Home	Ver Jun 8 2022
	Ring Doorbell	Up to date (11/2/2021)	Ring	3.45.1
Speaker	Ring Doorbell (*21, Wired)	Up to date (6/13/22)	Ring	Ver Jun 8 2022
	Echo Dot (4th Gen)	5805755780 (Software version)	Amazon Alexa†	2.2.432925.0
	Echo Show 5*	5805754756 (Software version)	Amazon Alexa†	2.2.432925.0
	Home Mini	Up to date (11/10/21)	Google Home†	2.45.1.8
	Nest Mini	Up to date (11/10/21)	Google Home†	2.45.1.8
	Homepod	15.0 (Software version)	Home (iPhone)†	Not available
	Homepod Mini	15.0 (Software version)	Home (iPhone)†	Not available
Home Appliance	Nest Hub Max*	43.2.26.392523459 (Software Version) 1.56.265669 (Cast Firmware Version)	Google Home†	2.45.1.8
	Samsung Fridge*	TIZEN 6.0 (AFH-US-KTM-21-XXXXMU 20210813_055302)		2.45.1.8
TV	GE Microwave	Up to date (11/5/21)	SmartHQ	1.0.0.101.11
	LG TV*	06.00.25 (Software Version) [LG] webOS TV UJ7700 (TV Information)		
	Samsung TV*	1290 (Software version)		
	Sony TV*			
Health	Vizio TV*			
	Oxylink Oxygen Monitor	Up to date (11/11/21)	ViHealth	2.72.0
	Renpho Smart Scale	Up to date (11/8/21)	Renpho	3.11.3
	Withings BPM Connect	Up to date (11/2/21)	Withings Health Mate	5.6.4
	Withings Sleep	Up to date (11/3/21)	Withings Health Mate	5.6.4
Withings Thermo	Up to date (11/2/21)	Thermo	2.0.0	

Table 5: The 57 devices used in this study, with device firmware and app software versions. This table supplements Table 1. App names and versions are left blank in cases where the device did not necessitate the use of a companion app.

Table 5 provides device firmware and app software version numbers where available. Devices that did not provide accessible version information are annotated with their recording date to best approximate the version used, as we factory-reset all devices and freshly installed all apps prior to video recording. This table is a supplement to Table 1.

Device Exclusions as User-Disadvantageous Designs The devices in Table 6 were removed from the study due to reasons articulated in subsection 3.1 and subsection 3.4. The behaviors causing

these exclusions carry implications for user harm and poor outcomes, particularly considering that IoT devices may be resold after being used. Two notable cases were requirements for payment information and inability to factory-reset devices.

For the former, we found that mandatory payment information for account registration—which also prevented prior studies from investigating certain apps or services [26, 42]—can be a dark pattern. The Wink Smart Hub, for example, required credit card information for a paid subscription service when creating a mandatory account via the companion app. This device required app use for

Device Type	Device Name	Ecosystem	Companion App Name	App Dependency	Reason for Exclusion
Home Automation	SmartLife Remote	D-Link	Smart Life	Smart interactions	Setup Issues
	KEYCO Air		KeyCo Air	All interactions	Reset Issues
	D-Link Mov		mydylink Home	All interactions	Setup Issues
Camera	Microseven Camera		Microseven	All interactions	Reset Issues
	Wansview Camera		Wansview	Setup interactions	Other Issues
Smart Hub	Wink Hub 2		Wink	All interactions	Payment Required
	IKEA Hub 2		IKEA Home Smart	All interactions	Setup Issues
Doorbell	iCSee Doorbell		iCSee	All interactions	Setup Issues
Speaker	Echo Spot	Amazon	Amazon Alexa	Setup interactions only	Redundant Model
	Echo Dot (two devices)	Amazon	Amazon Alexa	Setup interactions only	Redundant Model
	Echo Dot 3 (three devices)	Amazon	Amazon Alexa	Setup interactions only	Redundant Model
	Echo Flex (two devices)	Amazon	Amazon Alexa	Setup interactions only	Redundant Model
	Echo Plus	Amazon	Amazon Alexa	Setup interactions only	Redundant Model
Home Appliance	iKettle		Smarter 3.0	Smart interactions	Other Setup Issues
	Behmor Brewer		Behmor	Smart interactions	Account Registration Issues
	BlueAir Purifier	Samsung	BlueAir	Smart interactions	Reset & Setup Issues
	Smart Washer		SmartThings	Smart interactions	Reset & Setup Issues
	Smart Dryer		SmartThings	Smart interactions	Reset Issues

Table 6: Devices excluded from our study. Reset issues denote problems with effectively factory resetting the device to a fresh state, while setup issues (including account registration problems and payment requirements) prevented us from interacting with newly-reset devices. Devices with both reset and setup issues initially appeared to be correctly wiped and allowed for fresh registration, but otherwise indicated signs of incomplete or faulty factory resets. Redundant models were excluded to reduce duplicate interactions within this study.

all interactions, thus obstructing us from even setting up the device prior to selecting a subscription model. Upon further inspection, we learned that the subscription requirement was a recent update; prior to July 2020, Wink Hubs were at least partially free to use [96]. Requiring a paid subscription to support data storage and cloud functions is not concerning by itself, but requiring payment information at registration, prior to any device interactions, may prevent consumers from making informed decisions on whether to opt in to a longer-term financial relationship with the manufacturer or not. As such, these issues risk user exposure to financial and autonomy harms. In such cases, darkness may depend on appropriate disclosure: up-front transparency on long-term financial requirements may reduce darkness and deception.

With respect to resetting devices, we followed user-manual instructions for factory-resetting devices, taking cues from documentation to indicate “successful” reset (one device, the Microseven Camera, lacked a reset button and could not be reset via methods from product documentation). However, following reset, some companion-app-dependent devices could not be set up with fresh accounts due to issues like failed QR code scans or unexpected offline status. Worse, other devices like the Smart Washer and BlueAir Purifier appeared successfully reset (and therefore, wiped of user data), but displayed pre-reset data or otherwise indicated reset failure when connected to a fresh account or when re-connected to the established lab account in follow-up tests. We were forced to omit over 10% of all available devices in our lab from our study due to problems with reset and setup, as detailed in the Appendix (Table 6). These issues highlight how bugs in IoT devices may leave users with hardware that is inoperable or cannot safely be resold without sharing personal data with strangers. These outcomes relate to dark patterns like the *roach motel* and *forced action/disclosure*

(users are trapped with poor options to escape, or forced to disclose information), as well as dark pattern categories like *subverting expectations*.

A.2 Supplementary Tables

These tables provide further detail and context to other figures in section 4.

Top Dark Patterns Tables. To supplement subsection 4.4, we include Table 7, which lists the top five most commonly found unique dark patterns that we documented across all devices per type. We observe that the most common dark patterns in IoT interactions tend to be similar across types, which helps explain why the lowest overlap we observe in Figure 16a is $J = 0.27$.

Manufacturer Tables. We present the top five most popular unique dark patterns per manufacturer in Table 8. In general, the devices produced by each manufacturer were relatively consistent in the sense that they tended to incorporate all or most of the most frequent patterns within their cohort.

Modality Tables. Table 9 supplements the modality-based analysis in subsection 4.3 and Figure 13 and presents the top dark patterns by number of devices the pattern appeared in, broken down by devices’ degree of app dependency for our interactions.

A.3 Supplementary Figures

Figures that help provide context for other analyses, but may not be as independently illuminative as other figures in the paper.

Dark Pattern	Total №	Dark Pattern	Total №	Dark Pattern	Total №
Popup nag	2	Visual preference	11	Acc't required	8
No account deletion	2	Acc't required	9	Visual preference	8
Notifs preselected	2	Preselection - general	8	No checkbox for ToS/PP	7
Nonperm opt-out	2	No explanation of perm.	7	Notifs preselected	7
Acc't required	1	Popup nag	6	Nonperm opt-out	5
(a) Home Appliances (N=2)		(b) Home Automation (N=13)		(c) Cameras (N=10)	
Dark Pattern	Total №	Dark Pattern	Total №	Dark Pattern	Total №
Preselection - general	3	No explanation of perm.	4	Visual preference	7
Notifs preselected	3	Acc't required	3	Nonperm opt-out	6
Acc't required	2	Visual preference	3	Feature detour	6
Email sub. preselected	2	Extra notif badge	2	Preselection - general	6
Feature detour	2	No account deletion	2	Nagging - general	5
(d) Smart Hubs (N=5)		(e) Health Devices (N=5)		(f) Speakers (N=7)	
Dark Pattern	Total №	Dark Pattern	Total №	Dark Pattern	Total №
No notif settings	3	Acc't required	4	Preselection - general	7
Preselection - general	3	No checkbox for ToS/PP	4	Visual preference	6
Nagging self-promo	2	Notifs preselected	4	Nonperm opt-out	5
Confusing text	2	Visual preference	3	Nagging - general	5
No account deletion	2	Priv settings preselected	3	No account deletion	5
(g) TVs (N=4)		(h) Doorbells (N=4)		(i) Media Devices (N=7)	

Table 7: Top five dark patterns per device category, grouped by unique occurrences across devices within each category.




Figure 17: Average total counts of each pattern across devices, with each pattern colored according to our context categories. This figure provides the total (frequency) count complement to Figure 3.

Dark Pattern	Total №	Dark Pattern	Total №	Dark Pattern	Total №
Visual preference	122	Hidden feature behaviors	8	Visual preference	9
Nagging self-promo	41	Preselection - general	6	Nagging self-promo	5
Nonperm opt-out	32	Popup nag	6	Nonperm opt-out	4
No checkbox for ToS/PP	28	Notifs preselected	5	Unprompted suggestions	4
Preselection - general	23	Extra notif badge	4	Preselection - general	4
(a) Amazon (N=9)		(b) Samsung (N=3)		(c) Apple (N=3)	

Dark Pattern	Total №	Dark Pattern	Total №	Dark Pattern	Total №
Visual preference	102	Native ads	5	Visual preference	5
Preselection - general	35	No explanation of perm.	4	Acc't required	2
Nonperm opt-out	34	Acc't required	3	Priv settings preselected	2
Notifs preselected	14	Extra notif badge	3	Email sub. preselected	2
Feature detour	11	Nagging - general	3	No account deletion	2
(d) Google (N=7)		(e) Withings Devices (N=3)		(f) Kasa Devices (N=2)	

Dark Pattern	Total №
Visual preference	8
Notifs preselected	3
Acc't required	2
No checkbox for ToS/PP	2
Email sub. preselected	2
(g) Arlo Devices (N=2)	

Table 8: Top five dark patterns sorted by total frequency of each pattern across all our devices for that manufacturer. We include less-represented manufacturers Kasa and Arlo in these tables for illustrative purposes in section 5, but do not include these small manufacturers in subsection 4.5 figures due to small sample size.

Dark Pattern	Total №	Dark Pattern	Total №	Dark Pattern	Total №	Dark Pattern	Total №
Visual preference	26	Visual preference	7	Popup nag	3	Preselection - general	10
Acc't required	25	Nonperm opt-out	6	No explanation of perm.	3	Visual preference	9
Notifs preselected	20	Feature detour	6	Acc't required	2	No account deletion	8
No explanation of perm.	18	Preselection - general	6	Notifs preselected	2	Nonperm opt-out	8
Preselection - general	18	Nagging - general	5	Nonperm opt-out	1	Nagging self-promo	7
No checkbox for ToS/PP	17	No account deletion	5	Nagging - general	1	Nagging - general	6
No account deletion	16	Nagging self-promo	5	Hidden info	1	Priv settings preselected	6
Nonperm opt-out	16	Unprompted suggestions	5	Visual preference	1	Feature detour	5
Extra notif badge	15	No checkbox for ToS/PP	4	Preselection - general	1	Extra notif badge	5
Priv settings preselected	12	Popup nag	4	No account deletion	1	No notif settings	5
(a) App Only (N = 35)		(b) Both, For Setup (N = 7)		(c) Both, Non-setup (N = 3)		(d) Device Only (N = 12)	

Table 9: Top ten dark patterns for IoT devices that (1) must be controlled by a companion app, (b) require an app during setup, (c) only require an app for optional functionality, and (d) never prompted us for an app at all. N = 45 devices mandate some form of app use.

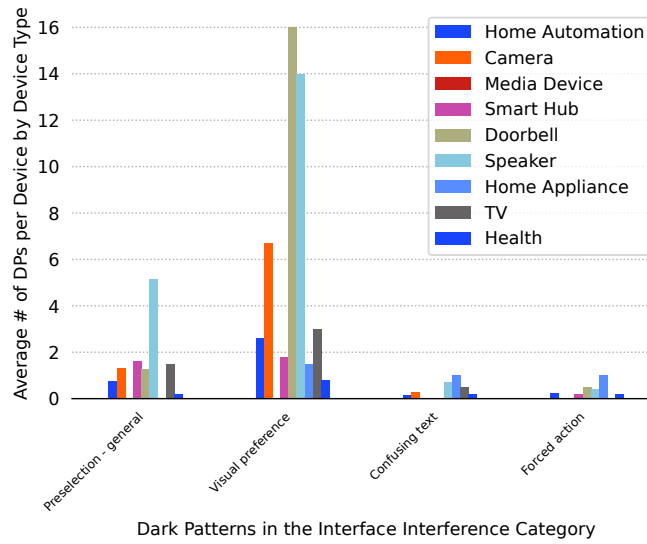


Figure 18: Average total counts of each pattern in the Interface Interference category per device type, with each pattern colored according to our context categories. This figure provides more context to the two highest average frequency patterns in Figure 3 and Figure 17, by inspecting Interface Interference patterns up-close and calculating the per-device-type average total counts (dividing the total count sum of each pattern per device type by the number of devices within that type, instead of the entire corpus of N=57).