

Evaluating the Contextual Integrity of Privacy Regulation: Parents’ IoT Toy Privacy Norms Versus COPPA

Noah Apthorpe

Princeton University

Sarah Varghese

Princeton University

Nick Feamster

Princeton University

Abstract

Increased concern about data privacy has prompted new and updated data protection regulations worldwide. However, there has been no rigorous way to test whether the practices mandated by these regulations actually align with the privacy norms of affected populations. Here, we demonstrate that surveys based on the theory of contextual integrity provide a quantifiable and scalable method for measuring the conformity of specific regulatory provisions to privacy norms. We apply this method to the U.S. Children’s Online Privacy Protection Act (COPPA), surveying 195 parents and providing the first data that COPPA’s mandates generally align with parents’ privacy expectations for Internet-connected “smart” children’s toys. Nevertheless, variations in the acceptability of data collection across specific smart toys, information types, parent ages, and other conditions emphasize the importance of detailed contextual factors to privacy norms, which may not be adequately captured by COPPA.

1 Introduction

Data privacy protections in the United States are enforced through a combination of state and federal legislation and regulatory action. In Europe, the General Data Protection Regulation (GDPR) is currently the best example of strong, centralized privacy legislation. The GDPR has inspired similar laws in other countries, such as the Brazilian General Data Privacy Law. According to the United Nations Conference on Trade and Development [51], 57% of countries have data protection and privacy legislation as of 2018.

Although data privacy protections vary across countries in terms of details and implementation, many share a common provenance: public pressure to protect sensitive personal data from unauthorized use or release. Surveys report that consumers worldwide were more concerned about online privacy in 2016 than 2014 [7] and that over 60% of U.S. survey respondents in 2018 are concerned about data privacy in general [34]. However, there has been no rigorous, quan-

tifiable, and scalable way to measure whether existing legal privacy protections actually match the privacy expectations of affected individuals. Without such data, it is difficult to know which aspects of privacy regulation effectively align company behaviors with social and cultural privacy norms and which necessitate further revision.

In this paper, we demonstrate that an existing survey technique [3] based on the formal privacy theory of contextual integrity (CI) [32] can be directly adapted to test the conformity of specific regulatory requirements to privacy norms, providing much-needed data to policymakers and the privacy research community. The survey technique can be applied to any privacy regulation that defines guidelines for data collection and transfer practices. Importantly, the survey technique involves questions describing privacy scenarios that are concrete and understandable to respondents from all backgrounds. It also allows straightforward longitudinal and cross-sector measurements to track the effectiveness of regulatory updates over time.

We present a rigorous case study of this technique evaluating the U.S. Children’s Online Privacy Protection Act (COPPA), which provides a federal legal framework to protect the online privacy of children under the age of 13. Specifically, we investigate whether parents’ opinions about the acceptability of data collection practices by Internet-connected “smart” children’s toys match COPPA mandates. Since the Federal Trade Commission (FTC) only updated its guidance on COPPA to explicitly include “connected toys or other Internet of Things devices” in June 2017 [16], our results provide the first indication as to whether COPPA aligns with parents’ privacy expectations.

This question is particularly relevant given the recent high-profile security breaches of smart toys, ranging from the theft of personal information of over 6 million children from toy manufacturer VTech to vulnerabilities in Mattel’s Hello Barbie [13]. More recently, Germany banned children’s smart watches and Genesis Toys’ My Friend Cayla doll, citing security risks and “spying concerns” [17, 31].

We survey a panel of 195 U.S. parents of children from

ages 3 to 13, the largest sample size for a study of parent opinions of smart toy data collection in the literature to date. We find that parents generally view information collection predicated on requirements specified by COPPA (e.g., “if the information is used to protect a child’s safety”) as acceptable, while viewing equivalent information collection without COPPA-specified conditions as unacceptable. This indicates that the existing conditions COPPA places on information collection by smart toys are generally in line with parents’ privacy norms, although there may be additional data collection requirements which could be added to regulation that were not tested in our study.

Additionally, we find that COPPA requirements for notification and consent result in more acceptable data collection practices than requirements related to confidentiality and security. This corroborates previous work indicating the primary importance of consent to user privacy norms [3]. We also find variations in the acceptability of COPPA-permitted data collection practices across specific smart toys, types of information, certain information use cases, parent ages, parent familiarity with COPPA, and whether parents own smart devices. These variations emphasize the importance of detailed contextual factors to parents’ privacy norms and motivate additional studies of populations with privacy norms that may be poorly represented by COPPA.

We conclude by noting that COPPA’s information collection criteria are broad enough to allow smart toy implementations that compromise children’s privacy while still adhering to the letter of the law. Continuing reports of smart toys violating COPPA [6] also suggest that many non-compliant toys remain available for purchase. Further improvements to both data privacy regulation and enforcement are still needed to keep pace with corporate practices, technological advancements, and privacy norms.

In summary, this paper makes the following contributions:

- Demonstrates that an existing survey method [3] based on contextual integrity [32] can be applied to test whether privacy regulations effectively match the norms of affected populations.
- Provides the first quantitative evidence that COPPA’s restrictions on smart toy data collection generally align with parents’ privacy expectations.
- Serves as a template for future work using contextual integrity surveys to analyze current or proposed privacy regulation for policy or systems design insights.

2 Background & Related Work

In this section, we place our work in the context of related research on contextual integrity, COPPA, and smart toys.

2.1 Contextual Integrity

The theory of contextual integrity (CI) provides a well-established framework for studying privacy norms and expectations [32]. Contextual integrity defines privacy as the appropriateness of information flows based on social or cultural norms in specific contexts. CI describes information flows using five parameters: (1) the subject of the information being transferred, (2) the sender of this information, (3) the attribute or type of information, (4) the recipient of the information, and (5) the transmission principle or condition imposed on the transfer of information from the sender to the recipient. For example, one might be comfortable with a search engine (*recipient*) collecting their (*subject & sender*) Internet browsing history (*attribute*) in order to improve search results (*transmission principle*), but not in order to improve advertisement targeting, which is a different transmission principle that places the information in a different context governed by different norms. Privacy norms can therefore be inferred from the reported appropriateness and acceptability of information flows with varying combinations of these five parameters.

Previous research has used CI to discover and analyze privacy norms in various contexts. In 2012, Winter used CI to design an interview study investigating Internet of things (IoT) device practices that could be viewed as privacy violations [54].

In 2016, Martin and Nissenbaum conducted a survey with vignette questions based on CI to understand discrepancies between people’s stated privacy values and their actions in online spaces [27]. Rather than straightforward contradictions, they find that these discrepancies are due to nuanced effects of contextual information informing real-world actions. This result motivates the use of CI in our study and others to investigate privacy norms in realistic situations.

In 2016, Shvartzshnaider et al. used the language of CI to survey crowdworkers’ privacy expectations regarding information flow in the education domain [46]. Survey respondents indicated whether information flows situated in clearly defined contexts violated acceptability norms. The results were converted into a logic specification language which could be used to verify privacy norm consistency and identify additional acceptable information flows.

In 2018, we designed a scalable survey method for discovering privacy norms using questions based on CI [3]. We applied the survey method to measure the acceptability of 3,840 information flows involving common connected devices for consumer homes. Results from 1,731 Amazon Mechanical Turk respondents informed recommendations for IoT device manufacturers, and regulators.

This paper adapts the survey method from our previous work [3] for a specific application: comparing privacy norms to privacy regulation. Our use of language from regulation in CI survey questions, direct comparison of discov-

ered privacy norms to policy compliance plans, and survey panel of special interest individuals (parents of children under age 13) distinguishes our work from previous uses of the survey method and previous CI studies in general.

2.2 COPPA & Smart Toys

Previous research has investigated Internet-connected toys and COPPA from various perspectives. Several studies have focused on identifying privacy and/or security vulnerabilities of specific smart toys [45, 48, 53], some of which are expressly noted as COPPA violations [6]. Our work uses these examples to inform the information flow descriptions included on our survey.

Researchers have also developed methods to automate the detection of COPPA violations. In 2017, Zimmeck et al. automatically analyzed 9,050 mobile application privacy policies and found that only 36% contained statements on user access, editing, and deletion rights required by COPPA [59]. In 2018, Reyes et al. automatically analyzed 5,855 Android applications designed for children and found that a majority potentially violated COPPA [43]. Most violations were due to collection of personally identifiable information or other identifiers via third-party software development kits (SDKs) used by the applications, often in violation of SDK terms of service. These widespread violations indicate that COPPA remains insufficiently enforced. Nevertheless, COPPA remains the primary legal foundation for state [30] and federal [12] action against IoT toy manufacturers and other technology companies for children’s privacy breaches.

Additional work has investigated parents’ and children’s relationships with Internet-connected toys. In 2015, Manches et al. conducted observational fieldwork of children playing with Internet-connect toys and held in-school workshops to investigate parents’ and children’s cognizance of how IoT toys work [26]. They found that most children and caregivers were unaware of IoT toys’ data collection potential, but quickly learned fundamental concepts of connected toy design when instructed.

In 2017, McReynolds et al. conducted interviews with parents and children to understand their mental models of and experience with Internet-connected toys [28]. Parents in this study were more aware of and concerned about IoT toy privacy than in [26], likely due to the intervening two years of negative publicity about connected toy privacy issues. The parents interviewed by McReynolds et al. provided feedback about desired privacy properties for connected toys, such as improved parental controls and recording indicators. The researchers urge ongoing enforcement of COPPA, but do not evaluate the parents’ responses in light of the law.

Our work builds on past research by obtaining opinions about smart toy information collection and transfer practices from a much larger pool of parents (195 subjects). We use these data to evaluate whether privacy protections mandated by COPPA align with parents’ privacy norms.

3 CI Survey Method

This study adapts a CI-based survey method first presented in our previous work [3] to evaluate whether specific requirements in privacy regulations align with user privacy norms. We chose this particular survey method because it is previously tested, scalable to large respondent populations, and easily adaptable to specific domains. The survey method works as follows, with our modifications for regulation analysis marked in *italics*:

1. Information transfers (“flows”) are defined according to CI as sets of five parameters: subject, sender, attribute, recipient, and transmission principle (described in Section 2.1).
2. We select lists of values for each of these parameters *drawn from or directly relevant to a particular piece of privacy regulation*. Using these values, we generate a combinatorial number of information flow descriptions *allowed or disallowed by the regulation*.
3. Survey respondents rate the acceptability of these information flows, each of which describe a concrete data collection scenario in an understandable context.
4. Comparing the average acceptability of flows *allowed or disallowed by the regulation* indicates how well they align with respondents’ privacy norms.
5. Variations in acceptability contingent upon specific information flow parameters or respondent demographics can reveal nuances in privacy norms *that may or may not be well served by the regulation*.

The following sections provide detailed descriptions of our survey design (Sections 3.1–3.2), deployment (Section 3.3), and results analysis (Section 3.4) for comparing parents’ privacy norms about smart toy data collection against COPPA regulation. Many of these steps mirror those in our previous work [3], but we include them here with specific details from this study for the sake of replicability.

3.1 Generating Smart Toy Information Flows

We first selected CI information flow parameters (Table 1) involving smart toys and specific data collection requirements from COPPA. We then programmatically generated information flow descriptions from all possible combinations of the selected CI parameters.

We next discarded certain information flow descriptions with unrealistic sender/attribute pairs, such as a toy speaker (sender) recording a child’s heart rate (attribute). Unrealistic sender/attribute pairs were identified at the authors’ discretion based on whether each toy could reasonably be expected to have access to each type of data during normal use. This decision was informed by smart toy products currently available on the market. The use of exclusions to remove

unrealistic information flows is a core part of the CI survey method [3] for reducing the total number of questions and the corresponding cost of running the survey. This process resulted in 1056 total information flow descriptions for use in CI survey questions (Section 3.2).

The degree to which these flows are rated as acceptable or unacceptable by survey respondents indicate agreement or disagreement between COPPA and parents' privacy norms. This rest of this section describes how we selected values for each information flow parameter in detail.

Transmission Principles from COPPA. We used the Federal Trade Commission's Six Step Compliance Plan for COPPA [10] to identify transmission principles. Some of these transmission principles match those in our previous work [3], facilitating results comparison.

We converted steps 2–4 of the Compliance Plan into four transmission principles regarding consent, notification, and privacy policy compliance (Table 1). COPPA dictates that parents must receive direct notice and provide verifiable consent before information about children is collected. Operators covered by COPPA must also post a privacy policy that describes what information will be collected and how it will be used. Our corresponding transmission principles allow us to test whether these requirements actually increase the acceptability of data collection from and about children.

The fifth step of the Compliance Plan concerns "parents' ongoing rights with respect to personal information collected from their kids" [10]. Operators must allow parents to review collected information, revoke their consent, or delete collected information. We translated this requirement into the transmission principle "if its owner can at any time revoke their consent, review or delete the information collected."

The sixth step of the Compliance Plan concerns operators' responsibility to implement "reasonable procedures to protect the security of kids' personal information" [10] and to only release children's information to third party service providers who can do likewise. We translated this step into five transmission principles involving confidentiality, security, storage and deletion practices (Table 1).

The Compliance Plan also lists a set of exclusions to COPPA. We converted the exclusions that were most applicable to Internet-connected children's devices into four transmission principles (Table 1). We also added the transmission principle "if it complies with the Children's Online Privacy Protection Rule" to test parents' trust and awareness of COPPA itself.

Importantly, we also included the *null* transmission principle to create control information flows with no COPPA-based criteria. Comparing the acceptability of flows with the *null* transmission principle against equivalent flows with COPPA-based transmission principles allows us to determine whether the COPPA conditions are relevant to parents' privacy norms.

Smart Toy Senders. The senders included in our survey represent five categories of children's IoT devices: a smart speaker/baby monitor, a smart watch, a toy walkie-talkie, a smart doll, and a toy robot. We chose these senders by searching for children's Internet-connected devices mentioned in recent press articles [13, 17, 20, 29, 31, 35], academic papers [5, 25], blogs [9, 19, 37], IoT-specific websites [21, 23, 38], and merchants such as Toys "R" Us and Amazon. All of the selected senders are devices that are reasonably "directed towards children" [10, 11] in order to ensure that they are covered by COPPA. We excluded devices such as smart thermometers or other smart home devices that might collect information about children but are not directly targeted at children.

It is important to note that the selected devices do not represent the full breadth of smart toy products. However, information flow descriptions involving specific devices or device categories evoke more richly varied privacy norms from survey respondents than flows describing a generic "smart toy." This is supported by existing interview data [58] noting that IoT device owners often have very different privacy opinions of specific entities than of their generic exemplars (e.g., the "Seattle government" versus "government").

Information Attributes. We reviewed academic research [25], online privacy websites [38], toy descriptions [15], and privacy policies [18, 36] to compile a list of information attributes collected by the toys in our sender list. The final selected attributes include heart rate, frequently asked questions, the times the subject is home, frequently traveled routes, the times the device is used, location, sleeping habits, call history, audio of the subject, emergency contacts, video of the subject, and birthday. These attributes cover a variety of personally identifiable or otherwise sensitive information with specific handling practices mandated by COPPA.

First- and Third-party Recipients. We included device manufacturers and third-party service providers as recipient parameters. This allowed us to examine variations in privacy between first and third parties while limiting the total number of information flows and the corresponding cost of running the survey.

Children as Information Subjects. The only subject parameter included in the survey is "its owner's child." This wording emphasizes that the child is not the owner of the device and acknowledges the parental role in ensuring children's privacy. It also accounts for devices that may not be used directly or exclusively by the child (e.g., a baby monitor). We indicated in the survey overview that respondents should think about their own children's information when interpreting this subject.

Sender	Transmission Principle
a smart speaker/baby monitor	<i>COPPA Compliance Plan Steps 2-3</i> if its privacy policy permits it if its owner is directly notified before the information was collected
a smart watch	
a toy walkie-talkie	
a smart doll	
a toy robot	
Recipient	
its manufacturer	
a third-party service provider	<i>COPPA Compliance Plan Step 4</i> if its owner has given verifiable consent if its owner has given verifiable consent before the information was collected
Subject & Attribute	
its owner's child's heart rate	<i>COPPA Compliance Plan Step 6</i> if it implements reasonable procedures to protect the information collected
its owner's child's frequently asked questions	if the information is kept confidential
the times its owner's child is home	if the information is kept secure
its owner's child's frequently traveled routes	if the information is stored for as long as is reasonably necessary for the purpose for which it was collected
the times it is used	if the information is deleted
its owner's child's location	
its owner's child's sleeping habits	
its owner's child's call history	
audio of its owner's child	<i>COPPA Exclusions</i> if the information is used to protect a child's safety
its owner's child's emergency contacts	if the information is used to provide support for internal operations of the device
video of its owner's child	if the information is used to maintain or analyze the function of the device
its owner's child's birthday	if the information is used to serve contextual ads
	<i>Other</i>
	if it complies with the Children's Online Privacy Protection Rule
	<i>null</i>

Table 1: Contextual integrity parameter values selected for information flow generation. The *null* transmission principle is an important control included to generate information flows with no explicit conditions. The transmission principles were derived from the FTC's Six Step Compliance Plan for COPPA [10].

3.2 Survey Design

We created and hosted the survey on the Qualtrics platform [39]. The survey was split into six sections: consent, demographic questions I, overview, contextual integrity questions, awareness questions, and demographic questions II. This section provides details about each section. The survey did not mention COPPA, privacy, security, nor any potential negative effects of smart toy information flows prior to the contextual integrity questions to prevent priming and framing effects.

Consent. Respondents were initially presented with a consent form approved by our university's Institutional Review Board. Respondents who did not consent to the form were not allowed to proceed with the study.

Demographic Questions I. The first set of demographic questions asked respondents for the ages of their children under 13. We chose this age limit because COPPA only applies to data collection from children under 13. We randomly

selected one of the ages for each respondent, n , which was piped to the survey overview.

Overview. Respondents were then presented with a survey overview containing a brief description of Internet-connected devices and instructions for the contextual integrity questions (Appendix A). This overview also explained how respondents should interpret the recurring phrase "its owner's child," and instructed them to keep their n -year-old child in mind while taking the survey (where n was selected for each respondent from their responses to the demographics questions I).

Contextual Integrity Questions. The core of the survey consisted of 32 blocks of questions querying the acceptability of our generated information flows (Section 3.1). Each question block contained 33 information flows with the same sender, same attribute, varying recipients, and varying transmission principles. For example, one block contained all information flows with the sender "a smart doll" and the at-

tribute “the times it is used.” Each question block also included one attention check question.

Each respondent was randomly assigned to a single question block. Answering questions about flows with the same sender and attribute reduced cognitive fatigue and ensured independence across recipients and transmission principles.

The information flows in each block were divided into matrices of individual Likert scale multiple choice questions. The first matrix in each block contained questions about information flows to different recipients with the *null* transmission principle (Figure 1). The remaining matrices each contained questions about information flows to a specific recipient with varying transmission principles (Figure 2). The order of the information flows in each block was randomized for each respondent.

Each individual multiple choice question in the matrices asked respondents to rate the acceptability of a single information flow on a scale of five Likert items: Completely Acceptable (2), Somewhat Acceptable (1), Neutral (0), Somewhat Unacceptable (-1), Completely Unacceptable (-2). We also included the option “Doesn’t Make Sense” to allow respondents to indicate if they didn’t understand the information flow.

Awareness Questions. Respondents then answered questions about their general technological familiarity and Internet use, ownership of Internet-connected devices, ownership of children’s Internet-connected devices, and previous knowledge of COPPA.

Demographic Questions II. Finally, respondents answered standard demographic questions from the United States Census. This allowed us to check the representativeness of our sample (Appendix B, Section 5.2) and account for demographic variables in our analysis.

3.3 Survey Deployment

We tested the survey on UserBob [52] once during the survey design process and again immediately prior to deployment. UserBob is a usability testing service for obtaining video screen capture of users interacting with a website while recording audio feedback. Each survey test involved creating a UserBob task with a link to the survey, brief instructions for users,¹ and settings to recruit 4 users to take the survey for 7 minutes each. UserBob automatically recruited users through Amazon Mechanical Turk at a cost of \$1 per user per minute. The resulting video and audio recordings of users interacting with the survey informed changes to our survey design. In particular, we reduced the number of questions per block and increased the number of pages over which the questions were presented. This reduced the amount of

¹UserBob task instructions: “This is a survey that will be given to a group of parents with children younger than 13. Take the survey, pretending you have one or more children younger than 13. Record your thoughts on the user interface and whether the questions do/don’t make sense.”

scrolling necessary to complete the survey and improved engagement. This practice of using pre-deployment “cognitive interviews” to test and debug survey design is common in survey research [49]. UserBob responses were not included in the final results.

We used Cint [8], an insights exchange platform, to deploy our survey to a panel of 296 adult parents of children under the age of 13 in the United States. We selected respondents with children younger than 13 because COPPA applies to “operators of websites or online services directed to children under 13” [11]. Our surveyed population therefore consisted entirely of individuals affected by COPPA. We chose not to set a minimum age for respondents’ children, because there is a lack of readily available information on the minimum age of use of Internet-connected children’s devices. While certain manufacturers list recommended minimum ages for their connected toys and devices, this was not the case for the majority of the devices we considered. Additionally, many devices such as wearable trackers, water bottles, baby monitors, are targeted towards very young children. Lastly, not restricting the minimum age allowed us to relax the demographic requirements for survey deployment.

Respondents were paid \$3 for valid responses where the attention check question was answered correctly. Each respondent was only allowed to answer the survey once. The survey responses were collected over an 18 hour time frame. We chose Cint to deploy our survey instead of Amazon Mechanical Turk, because Cint allowed us to directly target a specific panel of respondents (as in Zyskowski et al. [60]) without requiring a preliminary screening questionnaire to identify parents [44].

3.4 Response Analysis

We began with 296 responses. We removed the responses from 8 respondents who did not consent to the survey (none of their information was recorded) as well as those from 85 respondents who did not correctly answer the attention check question. We removed 2 responses in which over 50% of the information flows were characterized as “Doesn’t make sense.” We also removed 2 responses where not all information flow questions were answered. Finally, we removed 1 response where the respondent self-reported over 10 children and 3 responses that were completed in less than 2 minutes. This resulted in a final set of 195 responses with an average of 6 responses per information flow (standard deviation 1.4).

The responses to all contextual integrity questions (Section 3.2) were on a Likert scale with the following Likert items: “Completely acceptable” (2), “Somewhat acceptable” (1), “Neutral” (0), “Somewhat unacceptable” (-1), and “Completely unacceptable” (-2). We call this value the “acceptability score” of each information flow for each respondent.

In order to generalize privacy norms beyond individual respondents and information flows, we averaged the accept-

		Completely unacceptable	Somewhat unacceptable	Neutral	Somewhat acceptable	Completely acceptable	Doesn't make sense
its manufacturer		<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
a third-party service provider		<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 1: Example CI question matrix with information flows to different recipients and the *null* control transmission principle.

		Completely unacceptable	Somewhat unacceptable	Neutral	Somewhat acceptable	Completely acceptable	Doesn't make sense
if the information is stored for as long as is reasonably necessary for the purpose for which it was collected		<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
if the information is deleted		<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
if it implements reasonable procedures to protect the information collected		<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
if the information is kept confidential		<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
if the information is kept secure		<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 2: Example CI question matrix with information flows to a fixed recipient and varying transmission principles.

ability scores of flows grouped by CI parameters or respondent demographics. For example, we averaged the acceptability scores of all information flows with the recipient “its manufacturer” and the transmission principle “if the information is deleted” in order to quantify the pairwise effects of these two parameters on privacy norms. We then plotted these pairwise average acceptability scores as heatmaps to visualize how individual CI parameters or respondent demographic factors affect the overall alignment of information flows with privacy norms (Figures 3 & 4).

We statistically compared the effects of different COPPA provisions (Sections 4.1–4.4) by averaging the acceptability scores of all information flows grouped by transmission principles. For example, one group contained the average score given by each of the 195 respondents to information flows with non-*null* transmission principles, while a second group contained the average score given by each respondent to information flows with the *null* transmission principle. We then applied the Wilcoxon signed-rank test to find the likelihood that these two groups of scores come from the same

distribution. We performed three such tests with different transmission principle groups and set the threshold for significance to $p = 0.05/3 = 0.016$ to account for the Bonferroni multiple-testing correction.

We statistically compared the effects of smart device awareness, COPPA familiarity, and demographic factors (Sections 4.5–4.8) by averaging the acceptability scores of all information flows grouped by respondent category of interest. For example, one group contained the average score given by each respondent who owned a smart device across all answered CI questions, while the second set contained the average score given by each respondent who did not own a smart device. We then applied the Wilcoxon signed-rank test to find the likelihood that these two groups of scores come from the same distribution. We performed five such tests with groupings based on COPPA familiarity, age, smart device ownership, education, and income and set the threshold for significance to $p = 0.05/5 = 0.01$ to account for the Bonferroni multiple-testing correction.

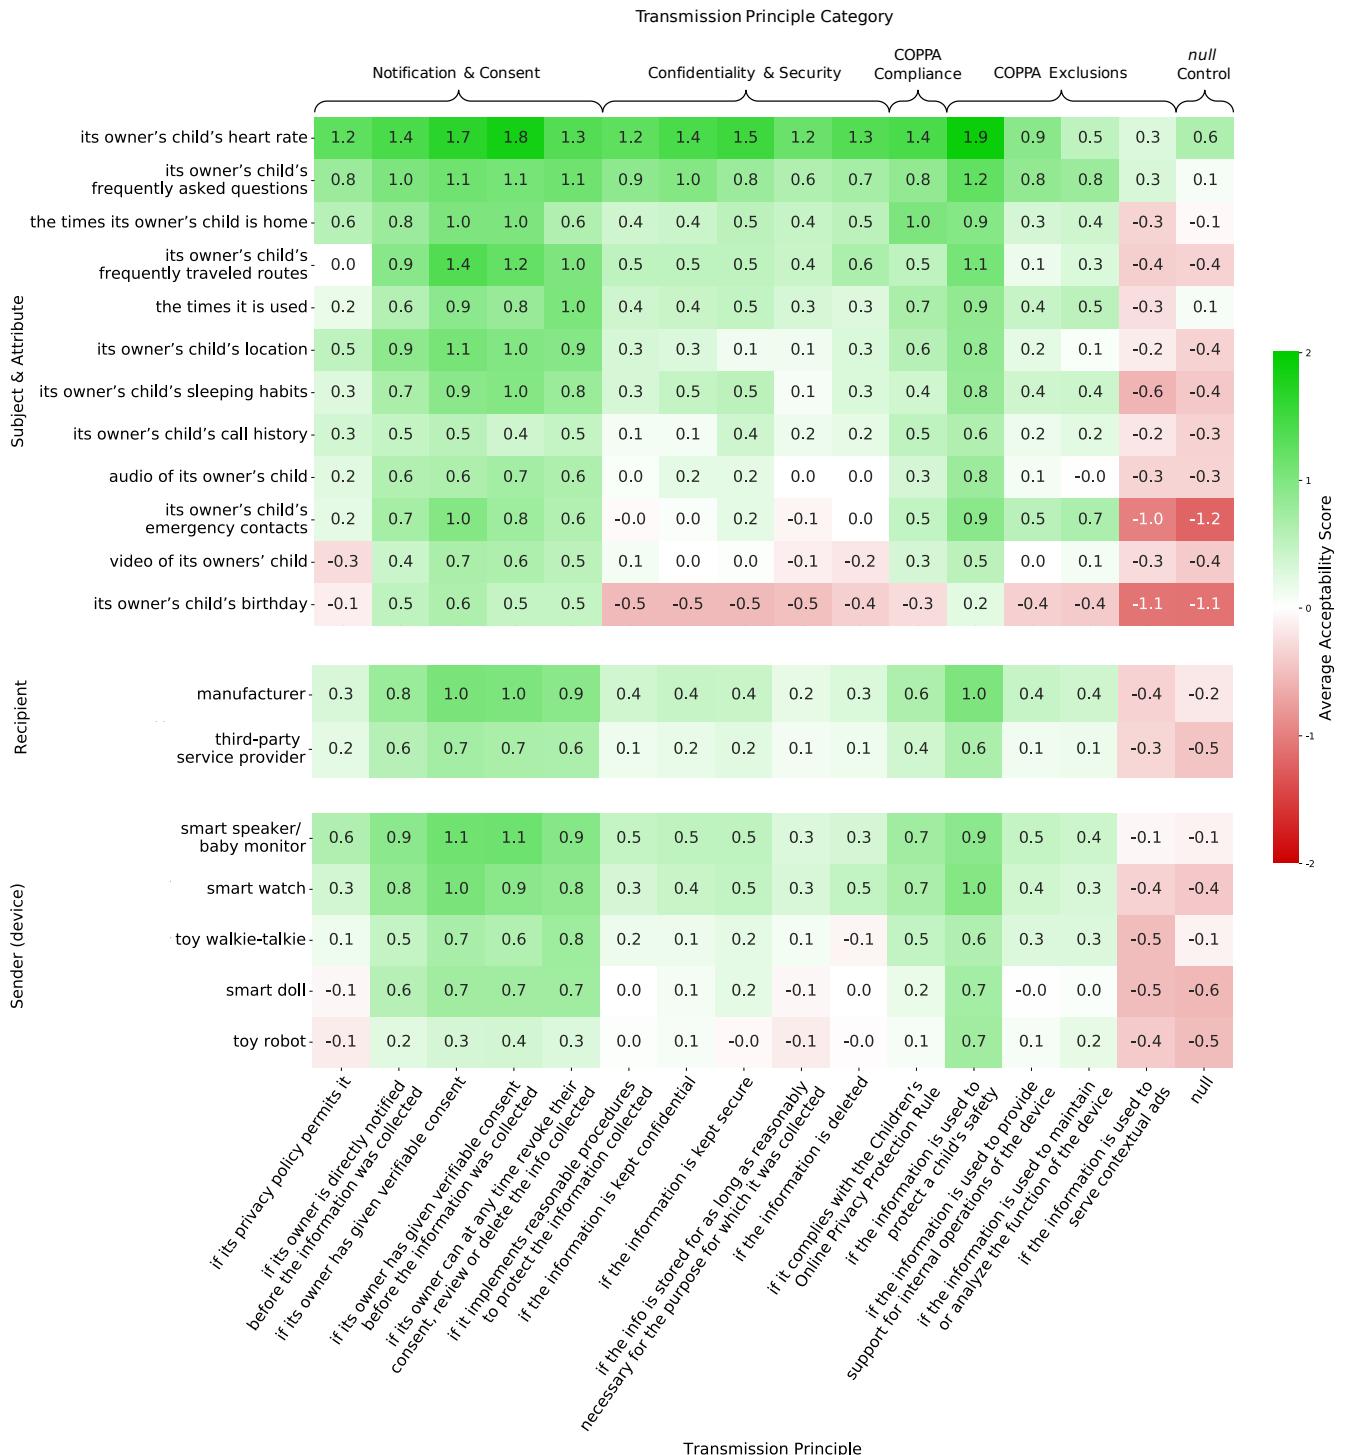


Figure 3: Average acceptability scores of information flows grouped by COPPA-derived transmission principles and attributes, recipients, or senders. Scores range from -2 (completely unacceptable) to 2 (completely acceptable).

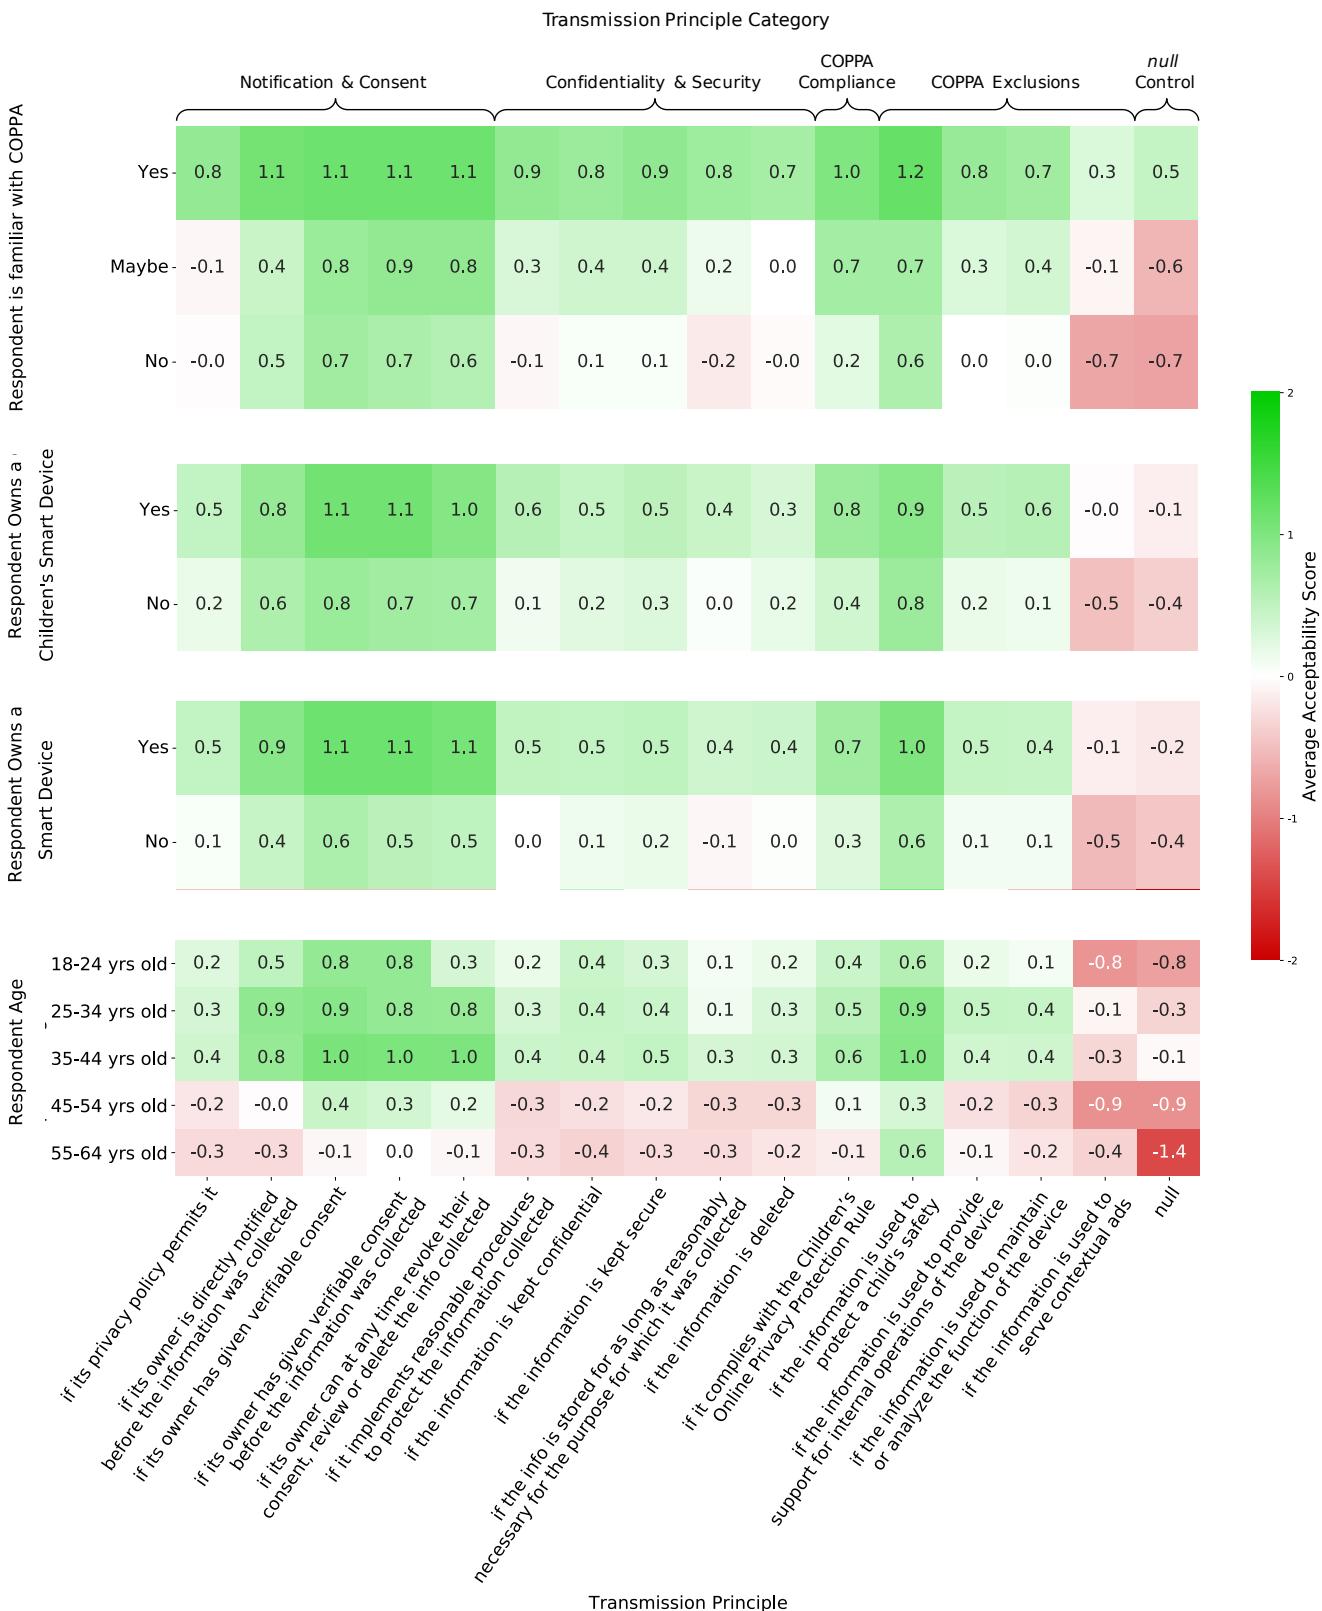


Figure 4: Average acceptability scores of information flows grouped by COPPA-derived transmission principles and respondent ages, familiarity with COPPA, or ownership of smart devices. Scores range from -2 (completely unacceptable) to 2 (completely acceptable)

4 Results

Overall, surveyed parents view information flows meeting COPPA data collection guidelines as acceptable while viewing equivalent information flows without COPPA criteria as unacceptable (Figures 3 & 4). This supports the conclusion that COPPA-mandated information handling practices generally align with parents' privacy norms. In this section, we elaborate on this finding and explore additional trends in our survey responses to further compare COPPA to parents' privacy norms regarding children's smart toys.

4.1 COPPA Data Collection Requirements Align with Parents' Privacy Norms

COPPA requirements were incorporated in the survey as information flow transmission principles derived from the FTC's Six-Step Compliance Plan for COPPA [10] (Section 3.1). The average acceptability scores of information flows explicitly obeying these requirements are mostly non-negative (Figures 3 & 4). This indicates that most surveyed parents consider these flows as "completely acceptable" or "somewhat acceptable." In comparison, the average acceptability scores of information flows with the control *null* transmission principle are mostly negative (Figures 3 & 4), indicating that most surveyed parents consider these flows without COPPA criteria as "completely unacceptable" or "somewhat unacceptable."

This difference between information flows with no explicit conditions versus flows with COPPA requirements holds regardless of information sender, recipient, attribute, or parents' demographics (apart from a few specific exceptions which we discuss below). On average, information flows with COPPA-derived transmission principles are 0.73 Likert-scale points more acceptable than their *null* transmission principle counterparts ($p < 0.001$).

Our research provides the first quantitative evidence that COPPA guidelines generally match parents' privacy norms for Internet-connected toys. This indicates that regulation can mandate meaningful transmission principles for information flows and supports further creation and fine-tuning of regulation to keep Internet data collection within the bounds of consumer privacy preferences.

4.2 Parents View Data Collection for Contextual Advertising as Unacceptable

Information flows with the transmission principle "if the information is used to serve contextual ads" have negative average acceptability scores across almost all senders, recipients, and attributes (Figure 3). Unlike all other information flows on our survey with non-*null* transmission principles, these flows are actually prohibited by COPPA. The "contextual ads" transmission principle is a "limited exception to COPPA's verifiable parental consent requirement" as

listed in the COPPA Compliance Plan [10]. This exception only applies to the collection of persistent identifiers (such as cookies, usernames, or user IDs) and not to any of the attributes included on our survey. Our respondents generally agree that collecting the attributes on our survey for contextual (targeted) advertising would be unacceptable, providing further support for COPPA's alignment with parents' norms.

This result indicates that the CI survey technique can detect regulatory provisions that reduce alignment with privacy norms, essential for future applications of the method (Section 6.2). It also provides evidence that the mere presence of a transmission principle doesn't necessarily improve the acceptability of information flows.

This result relates to existing work about opinions of data collection for advertising. Zheng et al. [58] interviewed owners of non-toy Internet-connected home devices and found mixed opinions of targeted advertising with data from these devices depending on the perceived benefit to the user. Combined with our results, this suggests that parents do not believe that relaxing COPPA to allow contextual advertising from more types of children's toy data would have enough benefit to outweigh privacy concerns.

4.3 Parents View Children's Birthdays as Especially Private

Information flows including the subject and attribute "its owner's child's birthday" are an exception to the trend described in Section 4.1. The average acceptability scores of information flows with this attribute and 10 of the 15 COPPA-derived transmission principles are negative (Figure 3). This discrepancy could be attributed to the relatively small number of parents (11 parents or 5.6% of total respondents) who were asked to score flows with this attribute. Parents may also view their children's birthdays as more personal than the other surveyed attributes or as less necessary for some of the surveyed transmission principles (such as "to maintain or analyze the function of the device"). Follow-up qualitative studies could focus on specific attributes, such as children's birthdays, to understand parents' rationales behind corresponding privacy norms.

4.4 Notification & Consent Versus Confidentiality & Security

Our results also provide insights into the relative importance of different sections within COPPA to parents' privacy norms. This could help regulators prioritize certain forms of non-compliant information collection for legal action.

Our COPPA-derived transmission principles can be divided into categories based on their topic and the section of the COPPA Compliance Plan [10] from which they were drawn (Section 3.1). One category consists of transmission principles from the Compliance Plan steps 2–5 regarding notification and consent (Table 1). These transmission principles

ples involve device privacy policies, the collection of verified consent, and the ability to revoke consent or review collected information. Another category consists of transmission principles from the Compliance Plan step 6 regarding information confidentiality and security (Table 1). These transmission principles involve reasonable data protection, confidential and secure storage, and limited information lifetime.

Across all senders, attributes, and recipients, information flows with transmission principles in the notification/consent category have significantly higher acceptability scores than flows with transmission principles in the confidentiality/security category by an average of 0.43 Likert scale points ($p < 0.001$) (Figure 3). One notable exception to this trend is the transmission principle “if its privacy policy permits it.” The acceptability scores for this transmission principle are an average of 0.53 Likert points lower than for others in the notification/consent category ($p < 0.001$). We suspect this reflects the general distrust of privacy policies evidenced in previous research [50]. Privacy policies are typically dense, lengthy, and difficult to interpret even for experts [42]. It therefore makes sense that parents would not view the disclosure of information collection in privacy policies as acceptable as other notification methods.

The greater acceptability of information flows with notification or consent criteria versus flows with confidentiality or security criteria corroborates previous research using the CI survey method to discover privacy norms of non-toy consumer IoT devices [3]. This provides longitudinal data indicating that users of Internet-connected products continue to prioritize consent over built-in security when reasoning about the appropriateness of information collection practices. This motivates continued work to improve the state of notification and consent mechanisms for Internet data collection. The most prevalent mechanisms, privacy policies and mobile application permissions, are widely understood to be ineffective for informing users or providing meaningful privacy control options [47]. As policies change to nuance the definitions of informed consent to include ideas of intelligibility, transparency and active opt-in, among others, it is important to continue to study and evaluate consumer’s privacy expectations regarding consent.

4.5 COPPA Compliance and Familiarity Increase Data Collection Acceptability

Information flows with the transmission principle “if it complies with the Children’s Online Privacy Protection Rule” received a positive average acceptability score of 0.49 across all senders, recipients, and attributes. As expected, flows with this transmission principle were rated as more acceptable by the 67% of respondents familiar with COPPA than by the 33% of respondents unfamiliar with the rule.

Furthermore, respondents who indicated that they were familiar with COPPA rated all information flows 0.75 Lik-

ert points more acceptable on average than respondents who were not familiar with the rule ($p < 0.001$) (Figure 4).

In both cases, stated compliance and/or familiarity with COPPA may increase parents’ acceptance of smart toy data collection by reassuring them that their children’s privacy is protected by regulation. However, this may be a false sense of security, as COPPA guidelines are relatively broad and COPPA violations are likely widespread in practice (Section 6.1) [6,43].

4.6 Younger Parents are More Accepting of Smart Toy Data Collection

Parents younger than 45 gave an average acceptability score of 0.48 to all rated flows, following the trend discussed in Section 4.1 (Figure 4). In comparison, parents 45 years and older gave an average acceptability score of -0.17 to all rated flows. This difference in the acceptability scores of these two groups is significant ($p < 0.01$). Nevertheless, context still matters, as information flows specifically “to protect a child’s safety” are viewed as generally acceptable to all surveyed parents regardless of age.

Previous work indicates that young American adults are more aware of online privacy risks and more likely to take steps to protect their privacy online than older adults [40]. Future studies could investigate why this awareness of online privacy risks makes younger parents more accepting of smart toy data collection.

4.7 Parents Who Own Smart Devices are More Accepting of Data Collection

Parents who own generic smart devices or children’s smart devices were more accepting of information flows than respondents who do not own these devices on average, but the difference in scores (0.34 Likert scale points) between these two groups is not significant ($p = 0.12$).

Nevertheless, this difference corroborates previous work using the CI survey method, in which owners of non-toy consumer IoT devices were found to be more accepting of information flows from these devices than non-owners [3]. This difference likely reflects a self-selection bias, in which those more uncomfortable with Internet data collection are less likely to purchase Internet-connected toys or other devices. However, the small effect size in both this study and the previous work may be due to parents purchasing smart toys unaware of their data collection potential [26] or willing to trade-off privacy concerns for other benefits provided by the products [58].

4.8 Education & Income have Little Effect on Parents’ Smart Toy Privacy Norms

Parents’ education and income did not have significant effects on acceptability scores. Parents earning more than

\$100,000 per year gave an average acceptability score of 0.46 to all rated flows, not significantly different from the average score of 0.37 from parents earning less ($p = 0.77$). Similarly, parents with at least some college education gave an average acceptability score of 0.37, not significantly different from the 0.33 average score of parents with a high school diploma or less ($p = 0.58$). This is perhaps unexpected given previous work indicating that parents with more resources are more likely to engage with children on privacy issues [41] and is a topic for follow-up research.

5 Limitations

Our results must be considered in the context of the following limitations.

5.1 Privacy Attitudes Versus Behaviors

Individuals often self-report greater privacy awareness and concerns than reflected in actual privacy-related behaviors [1, 22]. This “privacy paradox” is well-documented and poses a challenge for researchers. The CI survey method is vulnerable to privacy paradox effects. However, there is a reasonable argument that privacy regulation should prioritize the expressed norms of users (measured by the survey instrument) over norms evidenced through behaviors, which are influenced by external factors (such as confusing user interfaces) that could be affected by the regulation. For example, it is often difficult for consumers to determine the data collection practices of IoT devices, including Internet-connected children’s toys, due to poor company disclosure practices [42] and limited auditing by third parties. Just because many parents purchase smart toys does not mean that they approve of the toys’ data collection practices and wouldn’t support new regulation to improve privacy.

5.2 Respondent Representativeness

The self-reported demographics of our respondents (Appendix B) indicate that the sample, while diverse, is non-representative in ways that may influence measured privacy norms.

Females and high-income individuals are notably overrepresented in our sample compared to the United States population. The literature on gender differences in online privacy concerns suggests that women may generally perceive more privacy risks online than men [4, 14, 56], but some studies contradict this conclusion, reporting no significant gender effect [55]. The effect of income on online privacy concerns is similarly unsettled, with some reporting that high-income individuals are less concerned about privacy [24, 33], others reporting that high-income individuals are more likely to engage in privacy-preserving behaviors [41], and still others finding no significant income effect [57].

Limiting our survey to parents also ignores the opin-

ions of other parties, including school and daycare teachers and extended family members, who also purchase Internet-connected toys for children but may have different privacy norms. These individuals are also affected by COPPA and have legitimate justification for their opinions and interests to be reflected in children’s privacy regulation. Likewise, we did not ask whether our respondents were members of communities that may have less common privacy norms, but our respondent panel, drawn from across the United States, certainly missed smaller demographics.

Finally, our respondent panel consisted entirely of parents living in the United States, as COPPA only applies to products sold in the U.S. These respondents are therefore influenced by American attitudes toward privacy, which may vary from those of parents in other countries. We hope that future work will apply the CI survey method used in this paper to evaluate the alignment between privacy norms and privacy regulation in non-U.S. contexts.

5.3 Goals of Privacy Regulation

Our use of CI surveys to evaluate privacy regulation assumes that the underlying value of such regulation is to better align data collection practices with privacy norms. This makes an implicit normative argument about the purpose of privacy regulation, which does not necessarily hold, especially for the norms of majority populations. For example, privacy regulation may seek to protect minority or otherwise vulnerable populations. In these cases, surveys of all individuals affected by the regulation may reflect a majority view that does not value the norms or appreciate the situation of the target population. CI surveys could still be applied in these contexts, but care would need to be taken to identify and recruit respondents from populations differentially affected by the regulation in order to uncover discrepancies between the regulation and the norms of these groups.

Additionally, some regulation may be created with the goal of changing existing norms. In these cases, the CI survey method will indicate that the regulation does not match current privacy expectations upon enactment. However, CI surveys would still be useful for conducting longitudinal measurements to track whether the regulation has the desired effect on privacy norms over time.

6 Discussion & Future Work

We would like this study to serve as a template for future work using contextual integrity to analyze current or pending privacy regulation for policy or systems design insights. This section discusses our COPPA findings and presents suggestions for future applications of our method by policymakers, device manufacturers, and researchers.

6.1 COPPA Insights & Concerns

Previous research indicates that parents actively manage the information about their children on social media platforms to avoid oversharing [2], and that owners of IoT home appliances view most data collection by these devices as inherently unacceptable [3]. We expected that these domains would overlap, resulting in skepticism of smart toy data collection that even the restrictions in COPPA could not ameliorate. Surprisingly, it seems that the COPPA criteria assuaged parents' privacy concerns on average.

While we are encouraged that COPPA generally aligns with parents' privacy expectations, we are also concerned that the existence of COPPA may give parents an unreasonable expectation that their children's data is protected, especially since parents familiar with COPPA were less critical of smart toy information flows. In fact, several online services and Internet-connected toys have been found to violate COPPA [6,43], and many more non-compliant toys are likely available for purchase. Additionally, the information collection guidelines in COPPA are relatively broad, leaving room for technical implementations that adhere to the letter of the law but still compromise children's privacy. This motivates continued work by regulators and researchers to identify toys that place children's privacy at risk, as well as healthy skepticism by parents before purchasing any particular toy.

As an additional policy insight, variations in information flow acceptability across recipients² corroborate previous work [58] indicating that privacy norms are deeply contingent on the perception of entities that collect online data. COPPA distinguishes between first- and third-parties, but does not further categorize data recipients. This increases the flexibility of the law, but raises the potential that some recipients, which may be viewed as completely unacceptable by privacy norms, could still legally get access to children's data. This suggests that incorporating a more contextual framing of entities could improve the ability of future regulation to prevent unwanted data collection practices.

6.2 Further Policy Analysis Applications

The CI survey method is not limited to COPPA. We would like to see the results of follow-up studies focusing on different regulation, such as the Health Insurance Portability and Accountability Act (HIPAA), the Family Educational Rights and Privacy Act (FERPA), the National Cybersecurity Protection Advancement Act, the European General Data Protection Regulation (GDPR), and others from the local to international level, to see if their requirements result in similarly acceptable information flows for members of their target populations. As most privacy regulation encompasses information transfer or exchange, the theory of contextual integrity is an appropriate framework for this research. Further

studies would also allow cross-regulatory analysis to find common factors that affect alignment with privacy norms.

The CI survey method could also be incorporated into the policymaking process. Policy formulation and resource allocation could be guided by surveying a wide-variety of information flows allowed under current regulation and identifying egregious or unexpected norm violations that require attention. Policymakers could test whether previous regulatory approaches will be applicable to new innovations by conducting surveys with CI parameters describing new technologies and existing regulation (e.g., smart toys and COPPA prior to the 2017 inclusion of IoT devices [16]). Policymakers could also perform A/B tests of policy drafts with different stipulations and/or language by conducting multiple parallel surveys with varying CI parameters. These and other use cases would improve quantitative rigor in data-driven policy development and facilitate the design of regulation responsive to the privacy norms of affected populations.

6.3 Systems Design Applications

The application of CI surveys to guide systems and product design is covered in detail in our previous work [3]. To summarize, device manufacturers can conduct CI surveys to determine whether information collection practices of devices or new features under development will violate consumer privacy norms. This allows modifications during the design process to prevent consumer backlash and public relations debacles.

Applying CI surveys to evaluate privacy regulation can also yield valuable insights for systems research and development. For example, learning that parents value the ability to revoke consent or delete information (Figure 3) motivates research into verifiable deletion of cloud data from IoT platforms. Such insights are especially relevant as neither privacy norms nor regulations are necessarily tied to technical systems feasibility. Discovering that a particular CI parameter value is crucial to privacy norm adherence could launch several research projects developing efficient implementations or correctness proofs. We expect future applications of the CI survey method will generate many such results.

7 Conclusion

Increased interest in data privacy has spurred new and updated regulation around the world. However, there are no widely accepted methods to determine whether this regulation actually aligns with the privacy preferences of those it seeks to protect. Here, we demonstrate that a previously developed survey technique [3] based on the formal theory of contextual integrity (CI) can be adapted to effectively measure whether data privacy regulation matches the norms of affected populations. We apply this methodology to test whether the Children's Online Privacy Protection Act's re-

²Information flows to first-party manufacturers have higher average acceptability scores than flows to third-party service providers (Figure 3).

strictions on data collection by Internet-connected “smart” toys align with parents’ norms. We survey 195 parents of children younger than 13 about the acceptability of 1056 smart toy information flows that describe concrete data collection scenarios with and without COPPA restrictions.

We find that information flows conditionally allowed by COPPA are generally viewed as acceptable by the surveyed parents, in contrast to identical flows without COPPA-mandated restrictions. These are the first data indicating the general alignment of COPPA to parents’ privacy norms for smart toys. However, variations in information flow acceptability across smart toys, information types, and respondent demographics emphasize the importance of detailed contextual factors to privacy norms and motivate further study.

COPPA is just one of many U.S. and international data privacy regulations. We hope that this work will serve as a template for others to adopt and repeat the CI survey method to study other legislation, allowing for a cross-sectional and longitudinal picture of the ongoing relationship between regulation and social privacy norms.

Acknowledgments

We thank Yan Shvartzshnaider and our survey respondents. This work was supported by the Accenture Fund of the School of Engineering and Applied Science at Princeton University.

References

- [1] ACQUISTI, A., BRANDIMARTE, L., AND LOEWENSTEIN, G. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- [2] AMMARI, T., KUMAR, P., LAMPE, C., AND SCHOENEBECK, S. Managing children’s online identities: How parents decide what to disclose about their children online. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), ACM, pp. 1895–1904.
- [3] APTHORPE, N., SHVARTZSHNAIDER, Y., MATHUR, A., REISMAN, D., AND FEAMSTER, N. Discovering smart home internet of things privacy norms using contextual integrity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 2 (July 2018), 59:1–59:23.
- [4] BARTEL SHEEHAN, K. An investigation of gender differences in on-line privacy concerns and resultant behaviors. *Journal of Interactive Marketing* 13, 4 (1999), 24–38.
- [5] CHAUDRON, S., DI GIOITA, R., GEMO, M., HOLLOWAY, D., MARSH, J., MASCHERONI, G., PETER, J., AND YAMADA-RICE, D. Kaleidoscope on the internet of toys: Safety, security, privacy and societal insights. Tech. rep., EU Science Hub, Feb 2017.
- [6] CHU, G., APTHORPE, N., AND FEAMSTER, N. Security and privacy analyses of internet of things children’s toys. *IEEE Internet of Things Journal* (2018).
- [7] CIGI-Ipsos global survey on internet security and trust. <https://www.cigionline.org/internet-survey-2016>, 2016. Centre for International Governance Innovation.
- [8] Cint. <https://www.cint.com/>, 2018.
- [9] Consumerist Archives: COPPA. <https://consumerist.com/tag/coppa/index.html>, 2018. Consumer Reports.
- [10] FEDERAL TRADE COMMISSION. Children’s Online Privacy Protection Rule: a six-step compliance plan for your business. <https://www.ftc.gov/tips-advice/business-center/guidance/childrens-online-privacy-protection-rule-six-step-compliance>, July 2017.
- [11] FEDERAL TRADE COMMISSION. Children’s Online Privacy Protection Rule (“COPPA”). <https://www.ftc.gov/enforcement/rules/rulemaking-regulatory-reform-proceedings/childrens-online-privacy-protection-rule>, Aug 2017.
- [12] FEDERAL TRADE COMMISSION. Electronic toy maker vtech settles FTC allegations that it violated children’s privacy law and the FTC Act. <https://www.ftc.gov/news-events/press-releases/2018/01/electronic-toy-maker-vtech-settles-ftc-allegations-it-violated>, Jan 2018.
- [13] FINKLE, J., AND WAGSTAFF, J. Vtech hack exposes id theft risk in connecting kids to internet. <https://www.reuters.com/article/us-vtech-cyberattack-kids-analysis/vtech-hack-exposes-id-theft-risk-in-connecting-kids-to-internet-idUSKBN0TP0FQ20151206>, Dec 2015. Thomson Reuters.
- [14] FOGEL, J., AND NEHMAD, E. Internet social network communities: Risk taking, trust, and privacy concerns. *Computers in human behavior* 25, 1 (2009), 153–160.
- [15] Gator kids smart watch. <http://gatorsmartwatch.com/index.php/kids-gps-watch-supply/>, 2018.
- [16] GRAY, S. Federal Trade Commission: COPPA Applies to Connected Toys. *Future of Privacy Forum* (June

2017). <https://fpf.org/2017/06/26/federal-trade-commission-coppa-applies-connected-toys/>.

[17] HEATER, B. Germany bans smartwatches for kids over spying concerns. *TechCrunch* (Nov 2017). <https://techcrunch.com/2017/11/17/germany-bans-smartwatches-for-kids-over-spying-concerns/>.

[18] Hello barbie privacy policy. <https://toytalk.com/hellobarbie/privacy/>, April 2017. PullString.

[19] IoTList – Discover the Internet of Things. <http://iotlist.co/tag/kids>, 2018.

[20] JOHNSTON, P. Toy-telligence the smart choice for children at Christmas. <https://www.reuters.com/article/us-britain-christmas/toy-telligence-the-smart-choice-for-children-at-christmas-idUSKBN1CH283>, Oct 2017. Thomson Reuters.

[21] Kidsafe seal program: Member list. <https://www.kidsafeseal.com/certifiedproducts.html>, 2018. Samet Privacy, LLC.

[22] KOKOLAKIS, S. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & security* 64 (2017), 122–134.

[23] LAUGHLIN, A. Smart toys - should you buy them? <https://www.which.co.uk/reviews/smart-toys/article/smart-toys-should-you-buy-them>. Which? Digital Blog.

[24] MADDEN, M. Privacy, security, and digital inequality. *Data & Society* (2017).

[25] MAHMOUD, M., HOSSEN, M. Z., BARAKAT, H., MANNAN, M., AND YOUSSEF, A. Towards a comprehensive analytical framework for smart toy privacy practices. In *International Workshop on Socio-Technical Aspects in Security and Trust (STAST)* (2017).

[26] MANCHES, A., DUNCAN, P., PLOWMAN, L., AND SABETI, S. Three questions about the internet of things and children. *TechTrends* 59, 1 (2015), 76–83.

[27] MARTIN, K., AND NISSENBAUM, H. Measuring privacy: an empirical test using context to expose confounding variables. *Colum. Sci. & Tech. L. Rev.* 18 (2016), 176.

[28] MCREYNOLDS, E., HUBBARD, S., LAU, T., SARAF, A., CAKMAK, M., AND ROESNER, F. Toys that listen: A study of parents, children, and internet-connected toys. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), ACM, pp. 5197–5207.

[29] MOON, A. FBI warns parents of privacy risks associated with internet-connected toys. <https://www.reuters.com/article/us-usa-toys-fbi/fbi-warns-parents-of-privacy-risks-associated-with-internet-connected-toys-idUSKBN1A22AW>, Jul 2017. Thomson Reuters.

[30] NEW YORK STATE ATTORNEY GENERAL'S PRESS OFFICE. A.G. Schneiderman announces \$100,000 settlement with TRUSTe over flawed privacy certification program for popular children's websites, 2017.

[31] NIENABER, M. Germany bans talking doll Cayla, citing security risk. <https://www.reuters.com/article/us-germany-cyber-dolls/germany-bans-talking-doll-cayla-citing-security-risk-idUSKBN15W20Q>, Feb 2017. Thomson Reuters.

[32] NISSENBAUM, H. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.

[33] O'NEIL, D. Analysis of internet users level of online privacy concerns. *Social Science Computer Review* 19, 1 (2001), 17–31.

[34] PAUL, K. Here's the no. 1 issue that consumers want corporate America to fix today. <https://www.marketwatch.com/story/heres-the-no-1-reason-people-dislike-us-companies-2018-11-13>, Nov 2018. MarketWatch.

[35] PEACHMAN, R. R. Mattel pulls aristotle childrens device after privacy concerns. *The New York Times* (Oct 2017).

[36] pi lab privacy policy. <http://www.edwintheduck.com/privacy-policy>, 2018. pi lab.

[37] POLK, R. Internet of things: When it comes to smart toys, it pays to shop smart. <https://www.internetsociety.org/blog/2017/11/comes-smart-toys-pays-shop-smart/>, November 2017.

[38] Privacy not included. <https://advocacy.mozilla.org/en-US/privacynotincluded/>, 2018. Mozilla.

[39] Qualtrics. <http://www.qualtrics.com/>, 2018.

[40] RAINIE, L. The state of privacy in post-Snowden America. *Pew Research Center FactTank* (September 2016).

[41] REDMILES, E. Net benefits: Digital inequities in social capital, privacy preservation, and digital parenting practices of U.S. social media users. *International AAAI Conference on Web and Social Media* (2018).

[42] REIDENBERG, J. R., BREAUX, T., CRANOR, L. F., FRENCH, B., GRANNIS, A., GRAVES, J. T., LIU, F., McDONALD, A., NORTON, T. B., AND RAMANATH, R. Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Tech. LJ* 30 (2015), 39.

[43] REYES, I., WIJESEKERA, P., REARDON, J., ON, A. E. B., RAZAGHPANAH, A., VALLINA-RODRIGUEZ, N., AND EGELMAN, S. "Won't somebody think of the children?" Examining COPPA compliance at scale. *Proceedings on Privacy Enhancing Technologies 2018*, 3 (2018), 63–83.

[44] SCHLEIDER, J. L., AND WEISZ, J. R. Using Mechanical Turk to study family processes and youth mental health: A test of feasibility. *Journal of Child and Family Studies* 24, 11 (2015), 3235–3246.

[45] SHASHA, S., MAHMOUD, M., MANNAN, M., AND YOUSSEF, A. Smart but unsafe: Experimental evaluation of security and privacy practices in smart toys. *arXiv preprint arXiv:1809.05556* (2018).

[46] SHVARTZSHNAIDER, Y., TONG, S., WIES, T., KIFT, P., NISSENBAUM, H., SUBRAMANIAN, L., AND MITTAL, P. Learning privacy expectations by crowdsourcing contextual informational norms. In *Proceedings of the 4th AAAI Conference on Human Computation and Crowdsourcing (HCOMP16)* (2016).

[47] SOLOVE, D. J. Introduction: Privacy self-management and the consent dilemma. *Harv. L. Rev.* 126 (2012), 1880.

[48] STREIFF, J., KENNY, O., DAS, S., LEETH, A., AND CAMP, L. J. Who's watching your child? exploring home security risks with smart toy bears. In *IEEE/ACM Third International Conference on Internet-of-Things Design and Implementation (IoTDI)* (2018), pp. 285–286.

[49] SUDMAN, S., BRADBURN, N. M., SCHWARZ, N., ET AL. *Thinking about answers: The application of cognitive processes to survey methodology*. Jossey-Bass, San Francisco, CA, 1996.

[50] TUROW, J. Privacy policies on children's websites: Do they play by the rules?

[51] UNITED NATIONS CONFERENCE ON TRADE AND DEVELOPMENT. Summary of Adoption of E-Commerce Legislation Worldwide, January 2018.

[52] Userbob: Usability testing. <https://userbob.com/>, 2018.

[53] VALENTE, J., AND CARDENAS, A. A. Security & privacy in smart toys. In *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy* (2017), ACM, pp. 19–24.

[54] WINTER, J. S. Privacy and the emerging internet of things: using the framework of contextual integrity to inform policy. In *Pacific telecommunication council conference proceedings* (2012), vol. 2012.

[55] YAO, M. Z., RICE, R. E., AND WALLIS, K. Predicting user concerns about online privacy. *Journal of the American Society for Information Science and Technology* 58, 5 (2007), 710–722.

[56] YOUN, S., AND HALL, K. Gender and online privacy among teens: Risk perception, privacy concerns, and protection behaviors. *Cyberpsychology & behavior* 11, 6 (2008), 763–765.

[57] ZHANG, Y., CHEN, J. Q., AND WEN, K.-W. Characteristics of internet users and their privacy concerns: A comparative study between China and the United States. *Journal of internet Commerce* 1, 2 (2002), 1–16.

[58] ZHENG, S., APTHORPE, N., CHETTY, M., AND FEAMSTER, N. User perceptions of smart home iot privacy. In *Proceedings of the ACM on Human-Computer Interaction, Computer-Supported Cooperative Work and Social Computing (CSCW)* (November 2018).

[59] ZIMMECK, S., WANG, Z., ZOU, L., IYENGAR, R., LIU, B., SCHAUB, F., WILSON, S., SADEH, N., BELLOVIN, S. M., AND REIDENBERG, J. Automated analysis of privacy requirements for mobile apps. In *24th Network & Distributed System Security Symposium (NDSS 2017)*, NDSS (2017).

[60] ZYSKOWSKI, K., MORRIS, M. R., BIGHAM, J. P., GRAY, M. L., AND KANE, S. K. Accessible crowd-work?: Understanding the value in and challenge of microtask employment for people with disabilities. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (2015), ACM, pp. 1682–1693.

Appendix A: Survey Overview

Survey Overview

"Smart" Internet-connected household devices are becoming increasingly popular. Such devices often have sensors that collect information about the people who own and use them. This information may be sent to a variety of recipients for a variety of reasons.

This survey contains questions about information flows from devices used directly or indirectly by children. You will be asked whether you think each information flow is acceptable.

The information flows include the phrase "its owner's child." Here, "it" refers to the device, "owner" refers to the child's parent or legal guardian who purchased the device, and "child" refers to the child who is using the device or about whom the device collects information.

Please keep your 9 year old child in mind while answering the questions.

Please answer each question as honestly as possible.

Survey overview shown to participants before contextual integrity information flow questions. Participants are asked to keep one child in mind when answering the survey questions. The age of this child (9 in above example) is selected randomly for each participant from the self-reported ages of each of their children younger than 13.

Appendix B: Self-Reported Demographics and Technical Background of Survey Respondents

	Metric	Sample		Metric	Sample
Gender			Age		
	Female	61%	18-24 years old	3%	
	Male	39%	25-34 years old	31%	
Other/Prefer not to disclose	-		35-44 years old	48%	
Education			45-54 years old	13%	
	9th, 10th, 11th, 12th - no diploma	1%	55-64 years old	4%	
	High school graduate	14%	65 years or older	<1%	
	Some college but no degree	22%	Has 1 child	33%	
	Associate degree in college - Vocational	6%	Has 2 children	45%	
	Associate degree in college - Academic	4%	Has 3 children	15%	
	Bachelor's degree	30%	Has 4 or more children	7%	
	Master's degree	16%	answers based on 0-3 yr old child	14%	
	Professional school degree	2%	answers based on 4-7 yr old child	36%	
	Doctorate degree	5%	answers based on 8-12 yr old child	50%	
Income			0-3 hours of internet use per day	15%	
	Less than \$25,000	13%	4-7 hours of internet use per day	45%	
	Between \$25,000 and \$50,000	22%	8-12 hours of internet use per day	25%	
	Between \$50,000 and \$75,000	21%	>12 hours of internet use per day	14%	
	Between \$75,000 and \$100,000	21%	Uses a personal computer	97%	
	Between \$100,000 and \$200,000	17%	Uses a smartphone	94%	
	More than \$200,000	4%	Uses a tablet device	78%	
	Prefer not to disclose	2%	Owns a smart device*	49%	
Race			Does not own a smart device	50%	
	Asian	7%	Unsure	<1%	
	Black or African American	11%	Owns a children's smart device**	33%	
	Native Hawaiian or Other Pacific Islander	1%	Does not own a children's smart device	66%	
	White	76%	Unsure	1%	
	White, American Indian or Alaska Native	<1%	Familiar with COPPA	63%	
	White, Asian	<1%	Not familiar with COPPA	33%	
	White, Black or African American	1%	Maybe familiar with COPPA	4%	
	Other	3%			
Hispanic					
	Hispanic	14%			
	Not Hispanic	85%			
	Prefer not to disclose	<1%			

* Question text: "Do you own any 'smart' (Internet-connected) devices or appliances besides a smartphone, tablet, laptop, or desktop computer?"

** Question text: "Do you own any 'smart' (Internet-connected) devices or appliances used directly or indirectly by children besides a smartphone, tablet, laptop, or desktop computer?"