Can apps play by the COPPA Rules?

Ilaria Liccardi†*
ilaria@csail.mit.edu

Monica Bulger‡
monica@oii.ox.ac.uk

Hal Abelson†
hal@mit.edu

Daniel J. Weitzner†
djweitzner@csail.mit.edu

Wendy Mackay*
wendy.mackay@lri.fr

†Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA

‡Oxford Internet Institute
University of Oxford
England, UK

*INRIA Saclay Île-de-France
Orsay, France

***Abstract*—We review current technical and social barriers to COPPA compliance for popular online services aimed at children. We show that complying with COPPA has proven difficult for developers, even when a genuine attempt was made. We investigate reasons for this lack of compliance and identify common causes: specifically, difficulties obtaining verifiable parental control as well as supply mechanisms for parents to understand, review, grant access and monitor collection of their children's personal data. Unless part of online services, mobile apps do not need to comply with COPPA.**

We identify 38,842 (out of 635,264) apps which are self-described (by their developers) as suitable for young users. Half of these apps have the ability to collect personal data and only 6% present a privacy policy. Parents often have little to no knowledge or understanding of what data is accessed. Due to Android's design they must grant all access regardless of permission type or need. Among the self-described apps we find different levels of content rating; these are not a reflection of the content of the app itself but rather the required access to personal data.

We present a design for a new framework aimed at helping mobile apps to comply with COPPA. This framework aims to simplify the process for developers by providing appropriate tools and mechanisms to help comply with the COPPA rules while presenting an easily understandable interface for parents to review, navigate, understand and then grant access to their children's personal data.

Keywords*—*Privacy, children, mobile apps, COPPA.

I. Introduction

Children nowadays are technically literate and use technology that is available to them for chat, play and to communicate with friends; this will happen regardless of whether it is suitable for their age group. Many young children and teenagers have managed to create accounts on Facebook, by either lying about their age or using their parents' accounts (with their permission [1]). Parents tend to allow children to use their accounts, so they can monitor their children's online activities [1]. However, while monitoring children's access to online resources has proven beneficial when guarding children from possible threats [2], it is not sufficient to guard their online privacy. Parents might be unaware or uninformed of the possible collection, access and usage of their children's personal data.

In an effort to safeguard children's online privacy, the Federal Trade Commission introduced the **C**hildren's **O**nline **P**rivacy **P**rotection **A**ct (COPPA) to improve mechanisms for parents to control the information collected, used, and disclosed about their children's online behaviors. The legislation stipulates that websites actively collecting information from children under the age of 13 must seek written parental consent. COPPA was enacted in 1998, put into effect in 2000, and underwent revisions[1] in 2002, 2005, 2011, 2012 and 2013.

Central to the issue of implementation is COPPA's requirement that "operators" of commercial websites and online services obtain verifiable parental consent prior to collecting information about a child's participation. Further, COPPA ensures parental rights to review and amend access to their children's personal information after initial consent is granted. Past research [3] analyzed 162 popular children's websites and found that a staggeringly low number of these sites (only 4) fully complied with the major components of the law. This lack of compliance might be due either to the difficulty associated in obtaining verifiable and legitimate parental permission, or to the difficulties in supplying both information and mechanisms for parents to consent, review and monitor their children's personal data.

Many websites have enacted age-based bans to comply with COPPA, barring users under the age of 13 in their Terms of Service. Those which offer content to children and attempt to comply with COPPA guidelines follow a model similar to that of Disney's popular Club Penguin app[2]. When creating an account, a user under the age of 13 must provide an email address for their parent who is then sent an activation email. However, the identity of an email user is difficult to confirm [4] and as one mother demonstrated [5], aliases are easy to create. Further, age verification systems face technical and social challenges as young users attempt to subvert them for access [6], [1], [7]. While various children's websites and online services have tried to comply with COPPA's requirements, it is unclear how effective these safeguards are in limiting the children's disclosure of information [8].

The FTC's 2013 revisions of COPPA reflect an awareness of children's increasing and ever-changing technology use. This modification broadens the definition of *personal information* to include persistent identifiers such as cookies,

[1] Federal Register Notices:http://www.ftc.gov/enforcement/rules/rulemaking-regulatory-reform-proceedings/childrens-online-privacy-protection-rule

[2] http://disneyprivacycenter.com/privacy-policy-translations/english/

978-1-4799-3503-1/14/$31.00 ©2014 IEEE

geolocation, photos, videos and audio recordings[3][9]. While COPPA represents an effort to safeguard children's digital privacy, many argue that it is difficult to implement [3], resulting in loopholes for children's access to digital content and a failure to protect their information security.

These revisions reflect the need to to keep up with the increasing use of technology by young children. Research [10] has shown that an increasing number of children are accessing mobile applications via smartphones and tablets at younger ages. Recent Pew Internet studies found that 68% of children aged 12-13 owned a cell phone [10] and 71% accessed the Internet via a mobile device (phone, tablet, or other mobile device) [11]. Within this age group, 66% reported downloading a mobile app. These apps could have the ability to collect demographic, personally identifiable, behavioral and location data, though the types and combinations of data collected and usage varies from app to app. Often, children and their parents are unaware of the extent of data collected or are confused about practices of disclosure to third parties [1].

Is it feasible for websites and smartphone apps to comply with COPPA in light of these technical and social challenges? To determine the technical and social barriers to COPPA compliance, we first conducted a literature review of behavioral studies into children's engagement with apps and websites. We combine this with a technical analysis in which we also examine the age requirements for these apps in the form of the app's content rating. Each app's content rating determines the minimum age requirement for users to use the app.

We examine what kind of personal information these apps can collect about children, and whether notices in the form of privacy policies are provided within each app's page. We also analyze how parental controls can be implemented within the current Google Play store and suggest a possible framework that could achieve all of COPPA's requirements with fewer difficulties to developers or operators and simpler descriptions to parents.

II. Related Research

A. Social and technical challenges of remote age verification

In 2008, an Internet Safety Technical Task Force was formed to evaluate the role of technology in addressing children's online safety. In a national consultation with Internet service providers, social networking sites, policymakers, technology developers, educators, academics and child safety and public policy advocates, they evaluated the efficacy of existing technologies in promoting and safeguarding children's online safety. As part of their review, they found that age verification systems were "appealing in concept, but challenged in terms of effectiveness" [7]. Citing that systems which rely on remote verification of identity have "potential for inaccuracies" they describe public records or third-party in-person identity verification as more reliable, however these options can be costly and contend with additional challenges for implementation.

Age verification systems that rely on remote verification of identity face practical challenges from users who intentionally subvert them. Perhaps the most widely discussed example is Facebook, which, in its Terms of Service, forbids users under the age of 13[4]. Yet Boyd, et al. [1] report that millions of youth under the age of 13 use Facebook and lie about their age during site registration. When asked about age restrictions for Facebook use, parents were less likely to mention COPPA or privacy concerns and more frequently focused on issues of content and interaction. Parents surveyed in 2011 seemed to believe the age restriction related to mature content and interactions rather than a legal safeguarding requirement [1]. This confusion around COPPA's purpose is important, as if parents mistakenly think that the purpose of the legislation is to protect children from viewing upsetting content or contact, they may misunderstand its significance for safeguarding personal information.

O'Neill [6] observes that efforts by children and their parents to circumvent COPPA's restrictions raise ethical issues and asks *"who . . . bears responsibility for children's welfare in this context?"*. Often, it seems to be the children themselves. A Pew survey [10] of teen mobile app users found over half of respondents aged 12-13 to be wary of sharing personal information. 56% of respondents said they decided not to install a smartphone or tablet app after they discovered they would have to share personal information in order to use it, with 27% reporting uninstalling an app due to privacy concerns and 46% reporting turning off location tracking [10]. Yet younger children may not yet demonstrate the developmental capacity to make decisions to safeguard their information.

B. Children's cognitive capacity for decision-making

Why is age 13 the delineator in COPPA legislation? Research shows that developmental stages influence how youth make decisions. In general, as children mature, they show increased sophistication when interacting with commercial materials [12]. Most studies find that prior to age 11, youth are largely uncritical or reliant on inappropriate criteria when determining trustworthiness of online content [13] [14] [15]. For example, in a study of 135 children aged 8-10 years, the presence of dynamic features were believed to reflect the trustworthiness of a website [16]. Children ranked a website with animated pictures of dogs as more trustworthy than a website containing the same text, but no pictures. In fact, children rated sites with advertising and no information about the site owner or author as highest in trustworthinesseven though they often believed the author of the page was the advertiser — indicating a potential misunderstanding of the relationship between content provider and advertiser.

Advanced technical skills and apparent experience with mobile apps may mask problematic methods of assessing the trustworthiness of apps and websites youth encounter. Studies of children's responses to advertising show how apparent understanding of one dimension may not indicate informed consent for another. Several studies show that starting at around 8 years old, children can identify a selling intent in advertising, such as commercials during television programs or advertising on websites [17], [18], [19]. In fact, the American Psychological Association and the United States Institute of Medicine have established 8 years as the minimum age at which children can cognitively understand the purpose of

[3] http://www.ftc.gov/news-events/press-releases/2013/07/revised-childrens-online-privacy-protection-rule-goes-effect

[4] Facebook (2013). Statement of Rights and Responsibilities. https://www.facebook.com/legal/terms

advertising [20] [21]. Yet, at this age, children fail to see a persuasive intent in advertising, such as bias or inflated claims [22]. In other words, they may identify a purpose without understanding the means used to achieve it, which could potentially result in ill-informed decisions about sharing their information.

As youth develop decision-making strategies, they demonstrate inconsistencies when describing choices they would make under different scenarios. In a phone survey of 304 10-17 year olds, Turow and Nir [23] found that youth appeared appropriately concerned about sharing personal information, with 79% believing teenagers should ask their parents before sharing personal information and 73% reporting that they review privacy policies before using a website. However, when the researchers presented a scenario in which a gift or prize was offered, 45% said they would share personal information in exchange for a cash or gift incentive. These decisions also depend on the type of information disclosed. Walrave and Heirman (2013) [24] discovered that teens aged 12-18 were more willing to disclose certain types of personal information, with a mean of 69.5% willing to shares details in their profile (first name, age, gender, hobbies, and favourite shops) compared with a mean of 22.2% willing to share their contact data (home address, home phone number).

When making choices to lie about their age, use a fake email address or otherwise subvert age verification systems, youth under the age of 13 may not completely understand consent. They may not comprehend that consenting to access to a game or social media website means consenting to disclosure of their personal information such as location, times of access, and friends being collected and shared.

C. Legal framework: Compliance and consequences

From a legal perspective, compliance with COPPA presents several challenges. For example, websites and apps targeted to adults may still be of interest to and used by children. In 2003, a group of consumer advocacy organizations alleged that Amazon.com was in violation of COPPA because children were allegedly able to post product reviews that included their full name and city and state of residence [25]. This case raised the issue of the extent to which a website might be considered to be "targeted" toward children and therefore be required to comply with COPPA.

For websites and apps more clearly targeted toward children, enforcement of COPPA is uneven, with a few landmark cases. Since COPPA's enactment in 2000, the FTC has brought 20 enforcement actions totaling over $7.6 million [26]. In 2006, the FTC fined Xanga.com, a social networking company who created games such as Farmville, $1 million for collecting information from users who registered birth dates stating that they were under 13 years of age [27]. In its largest penalty, the FTC fined Playdom, a company providing online gaming apps related to Disney and Marvel storylines, $3 million for collecting and sharing information about users under the age of 13 [26].

Further, broad studies of COPPA compliance yield mixed results. In 2001, Turow [28] identified 162 websites with a high percentage of child users and found that 10% violated COPPA by collecting information on their users but not providing a privacy policy link on the home page (a requirement to be compliant with COPPA). Research conducted in 2003 [3] found that of 162 popular children's websites, only four fully complied with the major components of COPPA. This lack of compliance might be due to the difficulty in obtaining verifiable and legitimate parental permission. In a survey by the FTC [29] of 400 mobile apps using the word "kids" in their product description, nearly 60% transmitted device ID information, yet only 20% mentioned this practice in their privacy disclosure. Even when privacy policies are provided, they are often, as Turow [28] noted, too complex to be easily read or, for parent's rights under COPPA, such as to review information collected about their children, to be understood.

Even apps popularly considered to be COPPA compliant do not fulfill all of the listed requirements. While privacy policies may be prominently displayed, they are often confusing to understand [3]. More importantly, while apps and websites may provide appropriate means for parents to consent, we could not find any that allowed parents to review the information collected about their child, which is a COPPA requirement. In Cai and Zhao's [3] extensive review of 117 websites identified by Nielsen as primarily used by children, they found only 16 to be fully COPPA compliant. However, COPPA compliance in their 2013 study addresses parental consent, but does not clarify whether the websites have mechanisms for parents to request disclosure of "the general kinds of personal information" collected about their children. [30].

III. Self-described children's apps

We collected metadata about 635,264 apps on the Google Play store[5] from March 2013 to May 2013. From these we identified apps that self-described as being suitable for children, either in the title or in the description of the app itself. Apps were identified by searching for words such as *kid*, *kids*, *child*, *children*, *preschooler* and *preschoolers*. If the app contained one or more of these terms, it was flagged as targeted at children. If a negation (*not*) was present within the same sentence where the word appeared, we did not flag the app. We found 38,842 apps self-described as targeted at young children (Table I). Of the apps targeted toward children, only 10.8% (4,202 apps) presented a privacy policy within their app page.

We also analyzed what permissions these apps requested; in particular we were interested in permissions that grant the ability to read personal data (ie. contacts, location, bookmarks etc.). However, the mere presence of these permissions does not imply any collection behavior, hence we only counted personal permissions where any possibility existed of disclosing (the presence of a Full Network permission) to third parties (eg. developers, advertisers etc.) [31]. While 50% of the apps (19,540 apps) did not request any personal data, the remainder requested varying levels of accesss to personal data. The reason for this access is not part of the permission information.

[5]We gathered different apps by performing searches for dictionary words on the Google Play website, and retrieving the page for each app that was found. The search results are split onto multiple pages, so we retrieved each page of search results; Google Play enforces a maximum limit of 20 pages of results for any given search. We used different dictionaries to collect the apps. A large English dictionary and dictionaries for French, Italian, and Spanish were combined to create different queries. The Google Play website enforces rate limiting if a large number of requests are made; we therefore included logic that would detect error messages, pause and retry.

TABLE I. APPS DESCRIBED BY THEIR DEVELOPERS AS TARGETED AT CHILDREN, GROUPED BY THE NUMBER OF SENSITIVE PERMISSIONS REQUESTED BY EACH APP, SHOWING TOTAL, FREE AND PAID SETS FOR ALL APPS AND FOR THOSE WHICH PRESENT A PRIVACY POLICY WITHIN THEIR APP'S PAGE.

NUMBER OF SENSITIVE PERMISSIONS	38,842 APPS TARGETED AT CHILDREN WITHIN THE APP'S DESCRIPTION			4,202 (of 38,842) APPS HAVE A PRIVACY POLICY IN THEIR PAGE		
	TOTAL	FREE	PAID	APPS	FREE	PAID
0	19,540	11,261	8,279	1,840	1,069	771
1	8,200	5,974	2,226	1,028	775	253
2	2,955	2,468	487	502	410	92
3	2,069	1,843	226	282	229	53
4	3,138	2,772	366	284	268	16
5	1,710	1,646	64	92	83	9
6	324	285	39	47	42	5
7	203	173	30	32	29	3
8	373	333	40	35	31	4
9	241	229	12	18	17	1
10	39	36	3	17	17	-
11	27	25	2	12	12	-
12	4	4	-	1	1	-
13	14	14	-	10	10	-
14	2	1	1	1	1	-
15	2	1	1	-	-	-
17	1	1	-	1	1	-

Personal data access does not relate to the presence of a privacy policy. Hence users who choose apps where privacy policies are absent could infer what kind of information the app has the ability to access from the phone. However, no information or explanation regarding the reasons, use or sharing of this information is provided. A privacy policy provides the option for users to attempt to understand how their personal information is collected, used, stored and shared. Yet most privacy policies are difficult to understand due to the jargon, legalese and vague terms, and are often therefore ignored by users. [32].

Apps are not required to have a privacy policy present – either as a link or text – within their pages. Having a privacy policy might even be damaging for companies if they contain mistakes or omissions (related to what information the app collected, how they are used, shared or stored), – either willingly or unwillingly. If the app's behavior does not reflect (and can be clearly proven) what its privacy policy specifies, the Federal Trade Commission can punish apps' developer for violating the terms of its stated privacy policy [33][6] [34].

Users can try to infer what an app does by looking at the different *permissions requested* although there is no clear way to differentiate between information required for functionality of the app and information used for other purposes (or both). Users can also refer to the *content rating*, used to rate the content of the app itself and to limit some features (ie. user generated content and peer-to-peer communication) and collection of location data for younger users. Developers publishing to the Google Play store have the responsibility to use Google's guidelines to assign a content rating to each app – in the form of *everyone*, *low maturity*, *medium maturity* and *high maturity* [7]. However, while a developer might describe an app as suitable for children within its description text or title, this does not necessarily mean that its content rating corresponds (eg., *everyone* or *low maturity*). In fact, high maturity content ratings are present in apps that have been self-described as appropriate for children (Table II).

The reason for this disparity between content rating and an app's self-description is due to the type of personal information accessed by the app. Certain features or access to personal information might affect the content ratings that developers select for the app. Apps that access, publish or share location

[6]http://www.ftc.gov/news-events/press-releases/2013/12/android-flashlight-app-developer-settles-ftc-charges-it-deceived

[7]https://support.google.com/googleplay/android-developer/answer/188189?hl=en

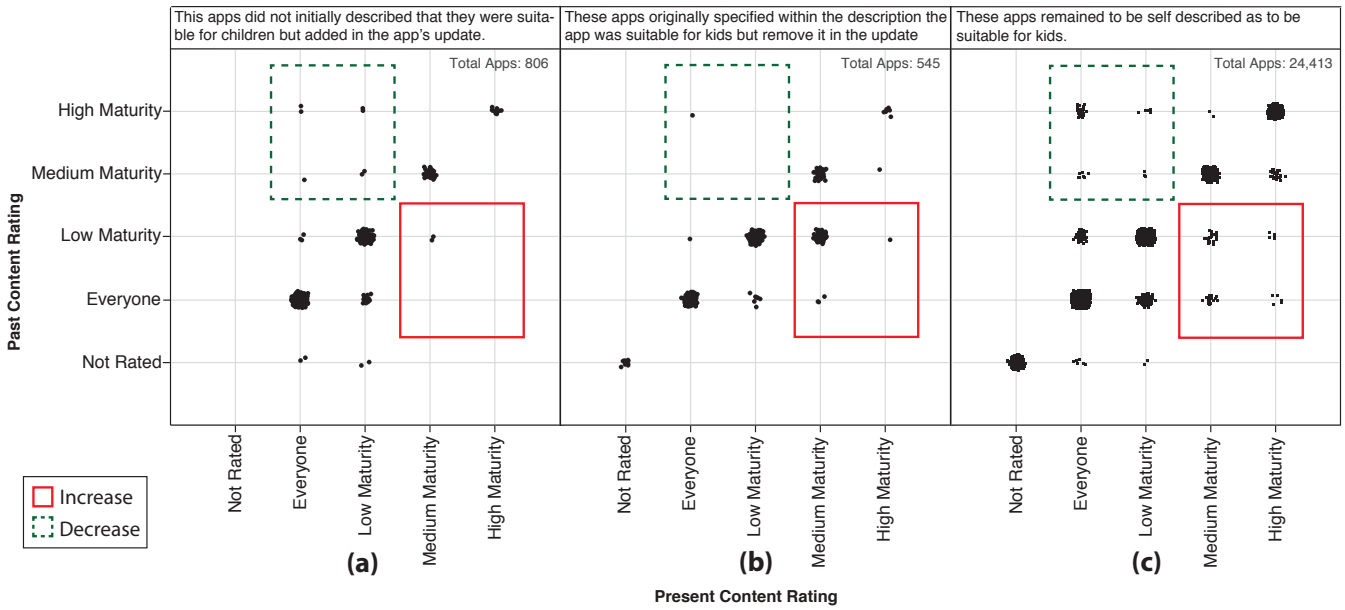

Fig. 1. Content rating can change between different versions of the same apps, either increasing the age suitability of the app or decreasing it. These changes are shown between three groups of apps: (a) apps that did not initially specify being suitable for kids but changed in their next version, specifying that they were suitable; (b) apps that initially described being suitable for kids but removed this in their next version; (c) apps that continued to assert that they were suitable for kids between the two versions of the app that were crawled.

TABLE II. CHILDREN'S APP CONTENT RATING

CONTENT RATING	TOTAL APPS	PRESENCE OF PRIVACY POLICY			
		WITHOUT		WITH	
		FREE	PAID	FREE	PAID
Everyone	24,610	13,924	8,183	1,534	969
Low Maturity	10,813	8220	1331	1065	197
Medium Maturity	1881	1126	403	323	29
High Maturity	1226	649	492	73	12
Not rated	312	152	160	0	0

data, host user generated content, or enable users to communicate and/or find each other, cannot be set in the lower maturity ratings (everyone, or lower maturity) even though the content of the app itself might be compliant with the content rating specification for the lower maturity rating. Developers wanting to access these types of personal data (for example location data for advertising purposes) and/or provide these types of functionalities will need to set the app's rating to a high maturity score in order to comply with Google rules.

After an initial maturity rating is established, developers can change, amend and add features and functionality to the app in newer versions. When new permissions are accessed, users are prompted, prior to updating the app, to accept these changes. Content ratings can also change between different app versions – a developer can increase or decrease the content rating, in some cases to reflect possible changes made to the app's functionality. However, changes in functionality are not required to change the content rating of an app. While permission information, if changed, prompts the user to accept new permissions, content rating changes are not flagged. Apps that are set to lower maturity ratings (*everyone* or *low maturity*) can change to higher maturity settings or vice-versa in new updates without users realizing that a change ever occurred (Figure 1)[8]

While Google Play offers apps that are tailored to younger users and apps with content ratings specifically designed to target children, it requires a user to have an account to be able to download an app. An account can only be created if the user themself to be older than 13. To comply with COPPA, the app store enacted age-based bans, banning users under the age of 13 in their terms of service. Children can also have educational accounts, in which case the responsibility for COPPA compliance shifts to the educational entity that has initiated the account. The reason for this approach is likely rooted in the fact that complying with COPPA has shown to be difficult to the point that even websites specifically aimed at children and designed to comply with COPPA's requirements do not, and perhaps cannot manage to do so [3].

To address the difficulty of complying with COPPA requirements, we developed a framework that can allow developers, companies and parents to be able to use, comply and understand all of COPPA's requirements and help safeguard the collection, usage, storage and sharing of children's personal information. This new infrastructure can help developers comply with COPPA requirements while alleviating technical challenges. Additionally, this framework can help parents understand commercial purposes for collecting, accessing and sharing their children's personal data.

[8]We collected app metadata from the Google Play store from October 2012 to January 2013 (preceding the version we are using in the paper), gathering information about 563,528 apps. We compared the two sets of apps and found 25,764 apps to be the same between the two datasets. We used the two versions of each app metadata and compared the assigned content ratings and self-descriptions as suitable for children

IV. COPPA-BASED INFRASTRUCTURE

Children's **O**nline **P**rivacy **P**rotection **A**ct (COPPA) addresses the handling of children's personal information in five key areas [9], [35]:

1) **Notices – Privacy Policies** *"Prominent and clear privacy notice must be provided on the home page and at each area where it collects personal information from children. The policy should contain who to contact, the kind of information collected, the usage and sharing intentions of the operators"* [9]. Apps complying with COPPA should have a link of their privacy policy within their own app's page for users to review it prior to installation. The privacy policy could also be presented during the permission request stage, or clear statements of purpose, usage and access could be embedded and tailored to each personal permission request.
2) **Consent Mechanism** *"Operators must obtain verifiable parental consent prior to collection of a child's personal data. An operator is required to send a new notice and request for consent to parents if there are material changes in the collection, use or disclosure practices to which the parent had previously agreed."* [9].

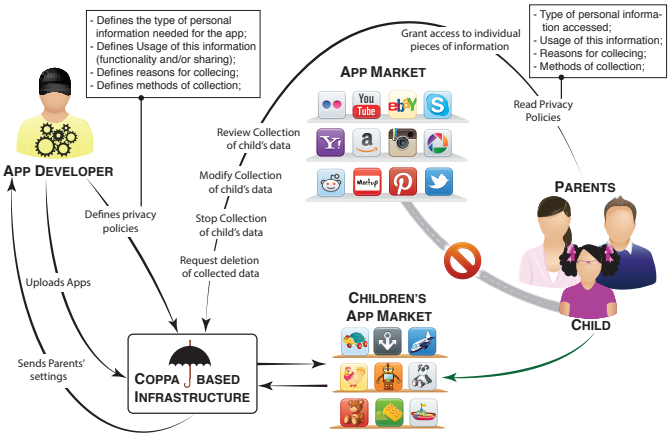

Fig. 2. Apps which choose to comply with COPPA must display and create a clear privacy policy. We show how our proposed infrastructure can help both parents and developers understand and communicate how personal information is collected, used and shared.

3) **Review the information** *"Operators must provide procedures whereby parents can review the child's personal information, request deletion or restrict future use, refuse to allow any further collection or use of the child's information. In turn operators that require certain personal information for functionality purpose might withdraw services if the information is not provided."* [9]. In apps, a clear distinction between information that is required for functionality and information that is required for additional purposes could be presented to parents prior to installation. Parents could decide to opt-out of additional purposes, while still allowing access for functionality.
4) **Limiting Collection** *"Operators may not require a child to disclose more information than is "reasonably necessary"*

to participate in an activity as a condition of participation when a child participates in online games and contests." [9]. Clear statements should be given to users as to why additional information is required. This information must be reviewed by parents prior to being disclosed.

5) **Confidentiality Security & Integrity** *"An operator must protect the confidentiality, security, and integrity of any personal data collected from children."* [9]. Apps' developers must specify how they are protecting users' data.

Our new infrastructure (Figure 2) helps both developers and parents to comply with, understand and review information based on each of these requirements. Parts of the tool will still require honest and clear descriptions provided by developers, however it also provides a mechanism for parents to request clarifications and flag possible misunderstandings and mistakes, to be either corrected by the developer of the app in question or resolved by the company responsible. More serious violations can be reported to the Federal Trade Commission.

In our framework, (Figure 2) children's accounts would be unique and access a limited subset of appropriate apps. In particular, children would not be able to browse any apps that do not comply with COPPA. A "virtual" separate app market would be browsable by children. Before children can use the app, parents must review and grant permissions using their own account. This app market[9] (and the associated parent's/developer's infrastructure) could be integrated as part of the current infrastructure (current app markets) or be completely separate, maintained and supported by a different entity (eg. a governmental one).

A. Notices - Privacy Policies

It might be difficult for developers to create a clear and concise privacy policy that describes how each app is functioning. Even when developers spend time creating clear and concise descriptions within privacy policies, parents might have difficulty comprehending technical and legal jargon. Additionally, since there are no stylistic guidelines or templates, each privacy policy could be uniquely organized and communicated. Hence a parent might need to re-read and understand different jargon within different policies.

In order to balance developers' and parents' needs we designed a tool to automatically generate a privacy policy

[9]While building this infrastructure is technically feasible and relative easy to implement with the right skill sets, the costs of hosting and monitoring such infrastructure are not investigated as part of the paper.

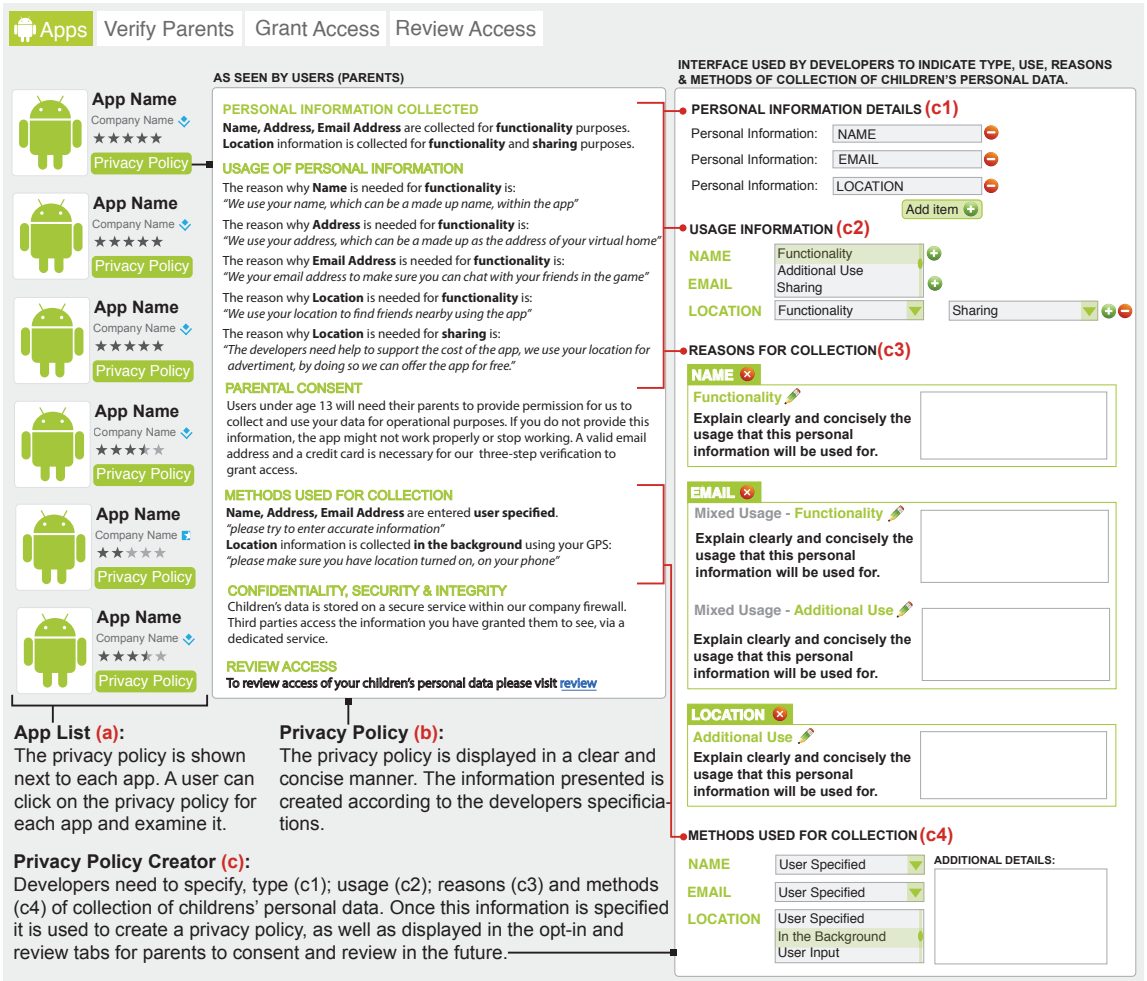

Fig. 3. Apps which choose to comply with COPPA must display and create a clear privacy policy. We show how our proposed infrastructure can help both parents and developers understand and communicate how personal information is collected, used and shared.

based on the developers' specifications of personal data used within the app. A pre-populated template will encourage concise and clear privacy policies for parents. If multiple apps use the same system, parents will gain familiarity with the system without needing to understand jargon specific to each app. The developer will need to provide information in four separate and interconnected steps (Figure 3 (c)):

1) **Personal Information details**: Developers need to specify all the personal information that they want to access. In this section the need to add the type of information (Figure 3 c(1)). For example if they want to access: name, email etc. Some of the personal information can be selected from a list of pre-defined terms, but the developers can add their own if necessary.
2) **Usage Information**: For each type of personal information, developers must specify the type of usage as either *functionality* (ie. required for the app to function), *additional use* (ie. used by the developer for some other use, like creating statistics) or *sharing* (ie. shared with third parties). If a particular piece of personal information is required for more than one usage, the developers must specify them in separate columns (Figure 3 c(2)).
3) **Reasons for collection**: Developers must specify in their own words reasons for the type of usage. When the same information has multiple usage types, all must be specified (Figure 3 c(3)). Developers write this section as free-form text. This section will be quoted within the auto-generated privacy policy.
4) **Methods used for collection**: Developers must specify the method by which the information is collected. They have choose between *user specified* ie. if the user has entered this information into the phone, *in the background* ie. if the information is collected silently in the background (for example location data) and *user input* ie. the user has to enter the information in a text box or select it from a list prior to starting the app (Figure 3 c(4)). In this section developers can also specify additional details in their own words.

After all information is entered, a privacy policy (Figure 3 (b)) is generated and displayed for parents to read. The privacy policy can be accessed by pressing the "privacy policy" button placed within the app's details (Figure 3 (a)).

B. Consent Mechanism

COPPA stipulates the provision of a means for parents to access their children's data, yet realizing this requirement presents both a technical and practical challenge. Proving via remote verification that an individual is the parent and not the child imitating them is difficult. A combination of factors might allow this to be more stable. A parental account could be an account that has been activated no sooner than 3 years ago[10] A parent will also supply a credit card, from which a $1 dollar is deducted with a description like "COPPA Protection" (this implements one of the options in the COPPA rules: 16 CFR 312.5(b)(ii)). In addition, a verification email will also be sent to the parents' account to be identified and authenticated. The proposed system triangulates communication and verification to make creating and using fake parental accounts more difficult. The system will help ensure developers are only in contact with parents who have been verified.

Of course, this system better addresses the technical than social challenges of identity verification. If a child can gain access to their parent's account, and/or their credit card and mail then they can likely spoof the system. This framework makes a concerted effort at identity verification, however we acknowledge the possibility that a parent may not realize their child is using their credit card and email accounts.

C. Review the Information

Reviewing and displaying collected information can be difficult for both developers and parents. Developers need to convey the information in a simple and concise matter for parents to understand. Parents need to be able to review the information and grant and remove relevant access. The system described in Figure 2 contains an easy-to-use visual interface that allows parents to grant access and later review the information that was collected about their children. If personal information is required for the app to function, the parents will need to manually grant access for each app's data collection (Figure 4 (a)). The parent must grant access to data that has been marked used for *functionality*. If parents want to see the purpose for each functionality, they can either hover over the information and see it displayed or view it in the privacy policy which can be displayed by pressing the *privacy policy* button (Figure 3 & Figure 4) next to the app's name. Parents can also browse the granted permissions by app (Figure 4 (a)), grouped according to the type of information (Figure 4 (b)) or by usage (Figure 4(c)).

D. Limiting Collection

This system does not have any way to actually control whether a developer is requiring more information than they need to. The parents can verify that each functionality request is in line with the data accessed. If this does not seem to be the case, the parent can flag the app for review (Figure 4 (a1)). The number of flags can be displayed within the app's page. This might incentivize developers to resolve the issue. The developer can either talk to the parents and resolve the issue. If this is not resolved and the flag is not removed, the market administrator can review and decide who is wrong. If developers are found to be at fault, the app in question can be removed from the market and a penalty might be inflicted on the company. In more serious cases, the FTC could also intervene.

E. Confidentiality, Security Integrity

It is the responsibility of the developers to safeguard the personally identifiable information of their users. Developers will need to explain in their own words in their privacy policy how this is achieved. It is a free form text that cannot be left blank and it is required to complete the processing of the privacy policy.

[10]We understand that for this to be the case, companies such as Google, Apple and Microsoft might need to provide a verification mechanism to query if the age requirement is met.

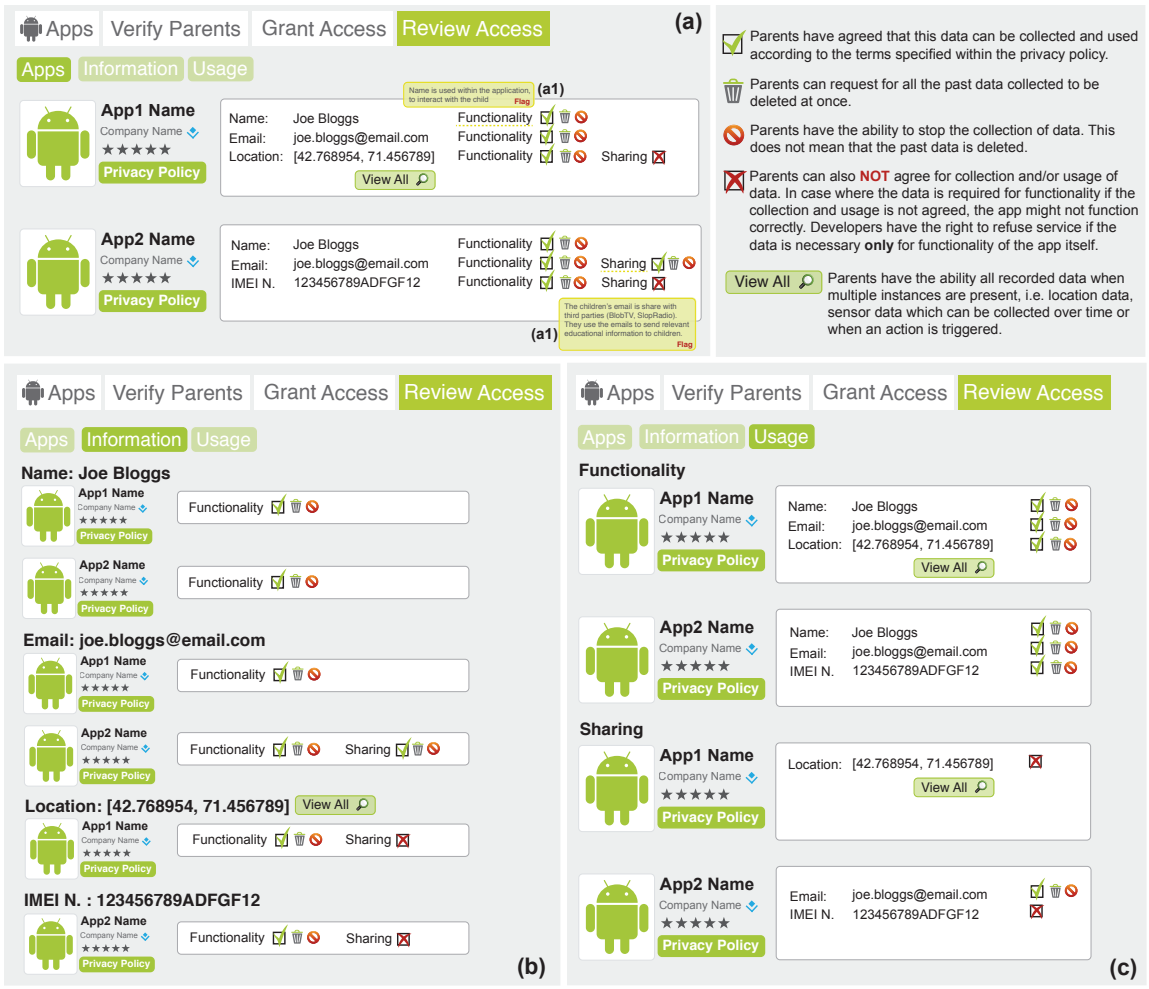

Fig. 4. Review of opt-in and opt-out of children's personal data collection showing usage and options for deletion, stopping and either revoking or not allowing access to the personal data. Parents can view this data, grouped by apps **(a)**, type of information **(b)** or usage of the data itself **(c)**.

V. Conclusion

We found that 50% of the apps we surveyed which were designated by their developers as appropriate for children had the ability to store and transmit children's personal information to third parties. However, the apps varied considerably in the types and amount of personal information they collected. Only 6% of these apps provided a privacy policy (Table I). A further consideration is that most privacy policies are difficult to understand due to length and jargon and are also visually difficult to read on a mobile device.

Our analyses revealed that mechanisms for parents to understand the suitability of an app may be misleading. In Google Play, parents may be misled if content ratings are relied upon that were selected by app developers according to Google's guidelines. Parents may believe that these ratings provide a full picture of an app's age-appropriate settings and content. Presumably, for example, the content ratings for this subset of apps would be set within the lower maturity ratings. However, our analyses show that this is not the case (Table II) and that there is a disparity between the actual content, based on the developers' descriptions, and the selected content rating of the apps. This disparity is due to the type of personal information accessed by the app. While the higher level of content maturity would seem to relate to violent, pornographic, or otherwise mature content, the ratings also reflect certain features or types of access to personal information. Apps cannot be set in the lower maturity ratings (everyone, or lower maturity) if they use location information, host user generated content, or enable users' communication, even if the content of the app itself is be compliant with the content rating specification for the lower maturity rating:

Developers wanting access to these types of personal data (for example location data for advertising purposes) or to provide functionality based on this kind of information, must set the app rating to a high maturity score in order to comply with Google's guidelines. However these maturity ratings primarily focus on the collection of location-related information. Developers are free to collect other types of personal information (eg. pictures, video, bookmark history) and still set the content to lower maturity ratings. Although software updates might result in changes to the content ratings (Figure 1), there are currently no clear mechanisms to inform parents (or users in general) of these changes.

Our combined literature review and technical analysis revealed both practical and technical challenges to complying with COPPA. In fact, apps in full compliance are rare due

to the many requirements of the legislation. While complying with COPPA can safeguard children's online privacy, research has shown that even when developers genuinely attempt to comply with the rules, that they fail to do so [3]. When developers have managed to comply, children find workarounds, such as spoofing systems to create fake parental accounts.

In this paper we have presented a new infrastructure which aims to help developers comply with the COPPA requirements while still allowing parents to simply and easily navigate and understand the access, usage, storage and sharing of their children's personal information. We have shown that this kind of system can be used to alleviate both the social and technical challenges and move toward realizing the core goals of COPPA, of safeguarding children's online access to their private information.

VI. Acknowledgment

Ilaria Liccardi was supported by the European Commission Marie Curie International Outgoing Fellowship grant 2011-301567 *Social Privacy*.

References

[1] d. boyd, E. Hargittai, J. Schultz, and J. Palfrey, "Why parents help their children lie to Facebook about age: Unintended consequences of the 'Childrens Online Privacy Protection Act'," *First Monday*, vol. 16, no. 11, 2011.

[2] M. Machill, "Internet Responsibility @ Schools: A Guideline for Teachers, Parents, and School Supervisory Bodies," *Kids On-line: Promoting Responsible Use and a Safe Environment on the Net in Asia, Shetty Kavitha ed. Singapore: Stamford Press, 227*, 2002.

[3] X. Cai, W. Gantz, N. Schwartz, and X. Wang, "Children's website adherence to the FTC's online privacy protection rule," *Journal of Applied Communication Research*, vol. 31, no. 4, pp. 346–362, 2003.

[4] E. J. Riegel, "Endangering Cyber Commerce," *Consumers Research Magazine, 81 (7), 345*, 1998.

[5] S. D. Estroff, "Undercover mom in clubpenguin, part 2: Let's get this party started!" *NetFamilyNews.org KID-TECH NEWS FOR PARENTS*, February 26 2012.

[6] B. ONeill, "Who cares? Practical ethics and the problem of underage users on social networking sites." *Ethics of Information Technology*, vol. 15, pp. 253–262, 2013.

[7] Internet Safety Technical Task Force, "Enhancing child safety online technologies: Final report of the internet safety technical task force," *The Berkman Center for Internet & Society at Harvard University*, 2008.

[8] M. O. Lwin, A. J. Stanaland, and A. D. Miyazaki, "Protecting children's privacy online: How parental mediation strategies affect website safeguard effectiveness," *Journal of Retailing*, vol. 84, no. 2, pp. 205 – 217, 2008.

[9] Federal Trade Commission, "16 CFR Part 312,," *Children's Online Privacy Protection Act of 1998, 5 U.S.C. 65016505, Bureau of Consumer Protection Business Center*, 2013, Final Rule.

[10] M. Madden, A. Lenhart, M. Duggan, S. Cortesi, and U. Gasser, "Tens and Technology 2013," *Pew's Internet and American Life Project & The Berkman Center for Internet & Society at Harvard University*, pp. 1–19, March 13 2013.

[11] M. Madden, A. Lenhart, S.Cortesi, and U. Gasser, "Teens and Mobile Apps Privacy," *Pew's Internet and American Life Project & The Berkman Center for Internet & Society at Harvard University*, pp. 1–20, August 22 2013.

[12] D. Roedder-John, "Consumer socialization of children: A retrospective look at twenty-five years of research." *Journal of Consumer Research*, no. 3, pp. 183–213, 1999.

[13] S. Bennett, K. Maton, and L. Kervin, "The 'digital natives' debate: A critical review of the evidence," *British Journal of Educational Technology*, vol. 5, pp. 775–786, 2008.

[14] E. Kuiper and M. Volman, "The web as a source of information for students in k-12 education." *In Coiro, J., Knobel, M., Lankshear, C. & Leu, D.J. (Eds.) Handbook of research on new literacies*, vol. 5, pp. 241–266. New York: Lawrence Erlbaum Associates, 2008.

[15] A. Walraven, S. Brand-Gruwel, and H. Boshuizen, "Fostering transfer of web searchers evaluation skills: A field test of two transfer theories." *Computers in Human Behavior*, vol. 26, pp. 716–728, 2010.

[16] M. Eastin, M.-S. Yang, and A. Nathanson, "Fostering transfer of web searchers evaluation skills: A field test of two transfer theories." *Journal of Broadcasting Electronic Media*, vol. 50(2), pp. 211–230, 2006.

[17] J. Harris, S. Speers, M. Schwartz, and K. Brownell, "US food company branded advergames on the Internet: Childrens exposure and effects on snack consumption," *Journal of Children and Media*, vol. 6, no. 1, pp. 51–68, 2011.

[18] V. Mallinckrodt and D. Mizerski, "The effects of playing an advergame on young childrens perceptions, preferences, and requests." *Journal of Advertising*, vol. 36, no. 2, pp. 87–100, 2007.

[19] K. Raine, T. Lobstein, J. Landon, M. P. Kent, S. Pellerin, T. Caulfield, D. Finegood, L. Mongeau, N. Neary, and J. Spence, "Restricting marketing to children: Consensus on policy interventions to address obesity," *Journal of Public Health Policy*, vol. 34, no. 2, pp. 239–253, 2013.

[20] National Research Council, "Food Marketing to Children and Youth: Threat or Opportunity," *Washington D.C.: The National Academies Press*, 2006.

[21] B. Wilcox, D. Kunkel, J. Cantor, P. Dowrick, S. Linn, and E. Palmer, "Report of the APA Task Force on Advertising and Children," *Washington D.C.: American Psychological Association.*, 2004.

[22] O. Carter, L. J. Patterson, R. Donovan, M. Ewing, and C. Roberts, "Childrens understanding of the selling versus persuasive intent of junk food advertising: Implications for regulation." *Social Science Medicine*, vol. 72, pp. 962–968, 2011.

[23] J. Turow and L. Nir, "The Internet and the family 2000. The view from parents. The view from kids," *Pennsylvania: The Annenberg Public Policy Center of the University of Pennsylvani*, 2000.

[24] M. Walraven and W. Heirman, "Adolescents, online marketing and privacy: Predicting adolescents willingness to disclose personal information for marketing purposes," *Children Society*, vol. 27, pp. 434–447, 2013.

[25] C. Burnes, "Data Protection," *Privacy and Data Protection*, vol. 3, no. 6, p. 12, 2003.

[26] J. Brill, "Address to family online safety institute," *FTC Public Statements*, p. 5, November. 15, 2012.

[27] J. ONeil, K. Drye, and C. Shannon, "Children's online privacy and protection act (COPPA) enforcement," *Privacy and Data Protection*, vol. 7, no. 1, p. 11, 2006.

[28] J. Turow, "Privacy policies on childrens websites: Do they play by the rules?" *Pennsylvania: The Annenberg Public Policy Center of the University of Pennsylvani*, 2001.

[29] Federal Trade Commission., "Mobile apps for kids: Disclosures still not making the grade," pp. 1–21, December 2012.

[30] Federal Trade Commission, "Mobile privacy disclosures: Building trust through transparency," pp. 1–29, February 2013.

[31] I. Liccardi, J. Pato, and D. J. Weitzner, "Improving User Choice Through Better Mobile Apps Transparency and Permissions Analysis," *Journal of Privacy and Confidentiality*, vol. 5, no. 2, Article 1, pp. 1–55, 2014.

[32] I. Liccardi, J. Pato, D. J. Weitzner, H. Abelson, and W. Mackay, "No Technical Understanding Required: Helping users make informed choices about mobile app access to their personal data." 2014.

[33] K. Tummarello, "Apple is paying back $32.5 million to parents of kids on app-buying sprees," *Bloomberg*, January 15, 2014.

[34] D. McLaughlin and A. Satariano, "Fandango, credit karma settle with ftc over app security flaws," *The Hill*, March 28, 2014.

[35] R. Malkin, "How the Children's Online Privacy Protection Act Affects Online Businesses and Consumers of Today and Tomorrow," *Loyola Consumer Law Review*, vol. 14, no. 2, 2002.