

Automating Accountability? Privacy Policies, Data Transparency, and the Third Party Problem¹

David Lie, Department of Electrical and Computer Engineering, University of Toronto

Lisa M. Austin, Faculty of Law, University of Toronto

Peter Yi Ping Sun, Bloomberg²

Wenjun Qiu, Department of Electrical and Computer Engineering, University of Toronto

1 Introduction

We have a data transparency problem. Currently one of the main mechanisms we have to understand data flows is through the self-reporting that organizations provide through privacy policies. However, it is notoriously difficult for individual consumers to read these policies and understand how their data is collected and used by the many organizations with whom they directly interact in our digitally mediated world. This problem is becoming more acute with the increasing complexity of the data ecosystem and the role of “third parties” -- the affiliates, partners, processors, ad agencies, analytic services, and data brokers involved in the contemporary data practices of organizations.³ These third party relationships were at the heart of the recent Cambridge Analytica scandal and are central to concerns about the surveillance capabilities of mobile devices.

There are many proposals to improve the usability of privacy policies, and transparency practices more generally, in relation to their role in enabling meaningful consumer consent.⁴ There are also important questions regarding whether privacy is best protected through such “self-management” paradigms.⁵ However, privacy policies disclose details of data flows and legal authority for processing that go beyond the question of consent obligations and raise the more general issue of accountability. Data transparency is important for ensuring accountability in data practices generally; without meaningful accountability we can have strong laws on the

¹ Our AppTrans project was funded by the Office of the Privacy Commissioner of Canada through their Contributions Program. We would also like to thank Robin Spillette, Mariana D’Angelo, and Michelle Wong for their excellent research assistance and input into the AppTrans project.

² The research Peter Sun performed for this paper was while he was a graduate student at the University of Toronto.

³ The distinctions between these can all be important in relation to some legal obligations but for the purposes of this paper we refer to them all as “third parties”. This is broader than how the term “third party” is defined in the GDPR, for example, and is more similar to how the GDPR defines “recipient”: See *GDPR*, *infra* note 9, art. 4. However, for much of this paper we are concerned with the practices of entities who would be considered third parties under the GDPR as well.

⁴ See s. 2 of this paper, below.

⁵ See, eg, Daniel J Solove, *Privacy Self-Management and the Consent Dilemma*, (2013) 126 Harv L Rev 1879.

books but weak compliance on the ground. It is reasonable to think that a space where nobody can detect wrongdoing and hold another to account is one where there is no effective legal accountability. There are several basic conditions for whether the data ecosystem is one that is actually governed by law. First, we need to be able to understand data flows. Second, we need to be able to check that these data flows are compliant with legal obligations. Third, we need to have sanctions for violations of the law. Privacy policies provide the possibility of fulfilling the first condition, but highly imperfectly and in a manner that makes it difficult to meaningfully integrate them with strategies to fulfill the second and third conditions. This makes broad data transparency elusive and highly dependent on the efforts of whistleblowers, investigative journalists, and civil society groups. Given the role of data in our society, and the broad range of rights and interests that it engages, this is not sustainable.

This paper addresses the question of whether we can improve the usability of privacy policies in relation to their role of enabling meaningful accountability. We argue that the automation of privacy policy analysis can make these policies more usable for at least some accountability purposes.

Our argument proceeds as follows. First, we outline why we need to shift from understanding privacy policies as a mechanism for informed consent and instead understand them as a transparency tool for enabling meaningful accountability. Although we conclude that automating privacy policy analysis has an important role to play in enhancing meaningful accountability in the data economy, we point to several limitations. These include the fact that privacy policies are essentially print technologies, even if digitized, and as such are “static” and have inherent limits in their ability to map onto dynamic data flows and properly disclose future data uses. Data transparency might also be better furthered through better technical methods for tracking data flows or for preventing them in the first place. Despite these limitations, we argue that automating privacy policies is an important step towards a better infrastructure of transparency and that policy changes to incentivize or require standardization of policies would help make that automation more accurate and effective.

Second, we discuss the literature on automating privacy policy analysis and show how existing projects focus on consumer-facing applications. As a contrast, we describe our AppTrans project, which involved prototyping a transparency tool for mobile app analysis that we developed to be regulator-facing. AppTrans combines automating privacy policy analysis with the auditing of data flows in order to flag apps that collect personal information without declaring this in their privacy policies. In the 700 apps we tested, we found a high rate of such discrepancies and also found that the main culprit for the discrepancies was data sharing with third parties such as the third party advertisers used to monetize the app. We describe the work we did to further refine the automated privacy policy analysis in order to provide a greater level of accuracy in relation to questions of third party sharing. We encountered significant difficulties in this work, and we analyze what this means in relation to the strategy of automating policy analysis in order to enhance transparency.

Third, in light of our findings we discuss whether such automated tools can help with the kind of proactive monitoring of policy compliance that several regulators have now called for in the context of the Facebook investigations in relation to the Cambridge Analytica scandal. Regulators in the US, Canada, and Europe have all pointed to the limitations of using

contractual safeguards to manage data sharing and have called for greater technical safeguards.

In conclusion, we argue that what we need is to shift from thinking about privacy policies as a transparency mechanism that enhances consumer choice and see them as a transparency mechanism that enhances meaningful accountability. We should view these policies as self-reported disclosures of data practices and treat them in a similar manner to other important forms of disclosures such as financial disclosures or tax reporting -- we need to regulate, standardize, and audit for compliance and we need to do it on a scale we have not even come close to managing so far in the privacy field.

2 Privacy Policies and Meaningful Accountability

2.1 The Limits of Focusing on Consumer Empowerment

Privacy policies are an important mechanism for data transparency but they are widely thought to be a failure in relation to improving consumer understanding of data flows. Most people do not read them, many find them difficult to understand, and even if people were to read and understand the policies directly relevant to the services they use this would take an unreasonable amount of time.⁶

Much of the focus on privacy policies is on whether consumers can understand their disclosures and this is because of the role of consent in private sector privacy regulation. For example, in the US, the FTC regulates privacy policies as an aspect of its jurisdiction over unfair and deceptive trade practices. Through this, the FTC has implemented a version of the Fair Information Practice Principles that, while more limited than in jurisdictions like the EU or Canada, still imposes strong requirements regarding notice and consent.⁷ As Solove and Hartzog argue, the FTC consent orders that have developed out of its complaint and settlement process have become the functional equivalent of a “common law” of privacy in the US. According to the FTC, insufficient notice includes “vague language tucked away in dense boilerplate agreements” and notice requirements are stricter where the personal information is “sensitive”.⁸

⁶ Aleecia M. McDonald & Lorrie Faith Cranor, “The Cost of Reading Privacy Policies” (2008) 4 I/S: J L & Pol’y For Info Soc’y 543.

⁷ Daniel J. Solove & Woodrow Hartzog, “The FTC and the New Common Law of Privacy” (2014) 114 Colum L Rev 583; For a comparison of the FTC to other FIPPs regimes, see Fred H. Cate, “The Failure of Fair Information Practice Principles” in Jane K. Winn, ed, *Consumer Protection in the Age of the Information Economy* (Abingdon: Taylor and Francis, 2006) 341.

⁸ Solove & Hartzog, *supra* note 7 at 635.

In Europe, the GDPR imposes very strict consent requirements. Consent must be opt-in rather than opt-out or implied, informed, freely given, and with the ongoing right to withdraw.⁹ However, under the GDPR, consent is one of six grounds for lawful processing so although consent requirements are strict there are other bases for grounding legal authority such as “legitimate interests” or “necessary for the performance of a contract”.¹⁰ Transparency requirements under such comprehensive data protection regimes are broader than what is required specifically for informed consent. Canada’s federal private sector legislation, PIPEDA, requires consent in more contexts than the GDPR, since it lacks some of the alternative authorities for data processing. However, in Canada consent can be implied or express, opt-in or opt-out, depending on the sensitivity of the information and reasonable expectations of users.¹¹ Because of the broad role of consent in Canadian privacy law, privacy policy disclosures are usually understood in relation to informed consent requirements. However, as in the case of the GDPR, transparency requirements in the law are actually broader than the disclosure requirements for informed consent.

Because of this general focus on the role of privacy policies in enabling meaningful consent and consumer choice, there has been a lot of important work done on making privacy policies more usable to consumers.

One set of strategies involves simplifying these policies, either in terms of their language or presentation, so that consumers can more easily understand them. For example, Canada’s proposed new federal privacy law, the *Consumer Privacy Protection Act*, imposes a “plain language” requirement.¹² An alternative to focusing on the language used is to focus on visualizations that aim at presenting the content of a privacy policy in a way that is quickly and easily accessible to the average person. For example, the FTC has recommended that app trade associations develop standardized icons and badges to enhance transparency.¹³ The Expandable Grid Project makes use of an interactive matrix with expandable rows and columns, and symbols.¹⁴ Kelley undertook a usability study and held that expandable grids may not actually be very user friendly and instead developed the nutrition label approach which attempts

⁹ EC, Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC, (General Data Protection Regulation), [2016] OJ, L 119/1, arts. 6, 4(11), 7 and recitals 32, 43 [hereinafter GDPR].

¹⁰ *Ibid* at art. 6.

¹¹ Personal Information Protection and Electronic Documents Act, SC 2000, c 5, Schedule One.

¹² Bill C-11, An Act to enact the Consumer Privacy Protection Act and the Personal Information and Data Protection Tribunal Act and to make consequential and related amendments to other Acts, 2nd Sess, 43rd Parl, 2020, c 62(1).

¹³ Fed. Trade Comm’n, “Mobile Privacy Disclosures: Building Trust Through Transparency” (February 2013), online (pdf):<https://www.ftc.gov/sites/default/files/documents/reports/mobile-privacy-disclosures-building-trust-through-transparency-federal-trade-commission-staff-report/130201mobileprivacyreport.pdf> at 25 ff.

¹⁴ Rob Reeder, “Policy authoring and Expandable grids”, online: *Robreeder.com* <http://www.robreeder.com/projects/xgrids.html>.

to display privacy policies in a single page summary.¹⁵ The format is inspired by the required format for food nutrition labels. Recently Apple became a high-profile adopter of this approach with the new privacy labels that it requires app developers to provide.¹⁶ Some of the problems with simplifying approaches is that they can oversimplify what are often complex dataflows, and they are unconnected to any process of verification of those dataflows so that consumers are not provided with assurances of their accuracy.

Another set of strategies involves enabling these policies to be machine readable. For example, P3P is a tool that defines a standard way of coding a policy in XML, a markup language for formatting text, which is easily machine readable.¹⁷ If a policy is machine readable then digital tools (referred to as “user agents”) can be created to enable a consumer to set their privacy preferences and automatically compare these to the privacy policies of a website that the consumer visits, flagging discrepancies for the consumer. In this way consumers do not themselves need to read each policy and so this strategy avoids the problem of using simplified language to map complex dataflows. Some P3P user agents combine strategies by adopting visualization techniques as a means of displaying P3P policies to users.¹⁸ Major drawbacks of the P3P approach include that it requires websites to make their policies machine readable with few incentives to do so, that it is often complex on the user side, and that there is no auditing or enforcement of P3P promises.¹⁹

More recently there has been widespread attention to the use of machine learning to automate privacy policy analysis in a manner that does not require coding a policy in a special language.²⁰ Rather than a human explicitly programming a machine to perform a task, machine

¹⁵ Patrick Gage Kelley et al., “A Nutrition Label for Privacy” (2009) Proceedings of the 5th Symposium on Usable Privacy and Security, online: <https://cups.cs.cmu.edu/soups/2009/proceedings/a4-kelley.pdf>.

¹⁶ Brian X. Chen, “What We Learned From Apple’s New Privacy Labels,” *The New York Times* (27 January 2021), online: <https://www.nytimes.com/2021/01/27/technology/personaltech/apple-privacy-labels.html>.

¹⁷ Lorrie Faith Cranor, “P3P: Making Privacy Policies More Useful” (2003) 1(6) *IEEE Security & Privacy* 50. In contrast, some standardization tools focus on helping create better privacy policies. See Carolyn A. Brodie, Clare-Marie Karat & John Karat, “An Empirical Study of Natural Language Parsing of Privacy Policy Rules Using the SPARCLE Policy Workbench” (2006) Proceedings of the Second Symposium on Usable Privacy and Security 8; D. Yu Weider and Savitha Murthy, “PPMLP: A Special Modeling Language Processor for Privacy Policies” (2007) 12th IEEE Symposium on Computers and Communications; Esmā Aimeur, Sebastien Gams, and Ai Ho, “Upp: User Privacy Policy for Social Networking Sites” (2009) Fourth International Conference on Internet and Web Applications and Services 267.

¹⁸ Robert W. Reeder et al., “A User Study of the Expandable Grid Applied to P3P Privacy Policy Visualization” (2008) Proceedings of the 7th ACM workshop on Privacy in the electronic society 45.

¹⁹ Electronic Privacy Information Center, “Pretty Poor Privacy: An Assessment of P3P and Internet Privacy” (June 2000), online: <https://epic.org/reports/prettypoorprivacy.html>.

²⁰ Rohan Ramanath et al., “Unsupervised Alignment of Privacy Policies using Hidden Markov Models” (2014) Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Short Papers) 605; Fei Liu et al., “A Step Towards Usable Privacy Policy: Automatic Alignment of Privacy Statements” (2014) 25th International Conference on Computational Linguistics 884; Rohan Ramanath et

learning offers a set of algorithms that can automatically create a program (sometimes called a model) that performs a task by observing (sometimes called “learning from”) many examples of input and output pairs that exemplify the task. For example, by observing many pictures of cats and things that are not cats, machine learning can train a model that can label whether a picture contains a cat or not. In this case, the data is privacy policies and machine learning techniques can be used to create models that can read and classify policies. This addresses the concern regarding incentives for adoption that affected P3P since all websites and mobile applications are already required by most jurisdictions to have a privacy policy. By utilizing already-existing policies, this strategy does not require organizations to do something new. Most of these projects are also consumer-focused in that they seek to automate privacy policy analysis in order to help consumers better understand privacy practices. Some of these projects also combine this automation with visualization methods. The Polisis project is an example of this, offering a machine learning tool for analyzing any privacy policy and then creating an interactive flowchart that allows individuals to visualize a summary of the policy.²¹ A number of other projects have involved crowdsourcing the interpretation of privacy policies in various ways in order to potentially help build tools to increase usability.²²

Automating privacy policy analysis, including automation that is paired with other tools such as visualizations, can potentially address some of the shortcomings of privacy policies in relation to consumer choice and can potentially do this better than simplification strategies that are deployed on their own, such as a focus on “plain language” or privacy labels. These shortcomings parallel the “cognitive” and “structural” problems that Solove usefully categorizes in his critique of the privacy “self-management” model.²³ Cognitive problems include the fact that policies are difficult to read and understand but also that people are not fully rational in their decision-making, leading individuals to make choices that are not consistent with their declared

al., “Identifying Relevant Text Fragments to Help Crowdsource Privacy Policy Annotations” (2014) Second AAAI Conference on Human Computation and Crowdsourcing 54; Noriko Tomuro, Steven Lytinen, and Kurt Hornsberg, “Automatic Summarization of Privacy Policies Using Ensemble Learning” (2016) Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy 133; Aaron K. Massey et al., “Automated Text Mining for Requirements Analysis of Policy Documents” (2013) 21st IEEE International Requirements Engineering Conference 4; Elisa Costante et al., “A Machine Learning Solution to Assess Privacy Policy Completeness: (short paper)” (2012) Proceedings of the 2012 ACM workshop on Privacy in the Electronic Society 91; Waleed Ammar et al., “Automatic Categorization of Privacy Policies: A Pilot Study” (2012) Carnegie Mellon Univ., online (pdf): <http://reports-archive.adm.cs.cmu.edu/anon/isr2012/CMU-ISR-12-114.pdf>; Sebastian Zimmeck & Steven M. Bellovin, “Privee: An Architecture for Automatically Analyzing Web Privacy Policies” (2014) Proceedings of the 23rd USENIX Security Symposium 1; Sebastian Zimmeck et al., “Automated Analysis of Privacy Requirements for Mobile Apps” (2017) 24th Network & Distributed System Security Symposium (2017).

²¹ See Hamza Harkous et al., “Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning” (2018), online (pdf): *PriboT*, https://pribot.org/files/Polisis_Technical_Report.pdf.

²² For a crowd-sourced approach instead, see Tos;DR, “Terms of Service; Didn’t Read” online: <https://tosdr.org/>; Shomir Wilson et al., “Crowdsourcing Annotations for Websites’ Privacy Policies: Can It Really Work?” (2016) Proceedings of the 25th International Conference on World Wide Web 133; Yan Shvartzshnaider et al., “Analyzing Privacy Policies Using Contextual Integrity Annotations” (2018) [unpublished], online: SSRN <https://ssrn.com/abstract=3244876>.

²³ Solove, *supra* note 5.

privacy preferences. Automating privacy policies is a good starting point for developing technological tools that would allow individuals to understand different aspects of the policies without necessarily having to read them, which is what the Polisis project aims at with its “AI-powered summary” of any privacy policy.²⁴ Automating policies could also aid in the creation of some tools that would allow individuals to indicate their privacy preferences in general and then have a tool check various policies to flag features inconsistent with these preferences.

Structural problems of the self-management model include problems of scale (too many organizations now collect information about individuals, making “self-management” difficult), aggregation (information aggregated over time raises potential privacy issues not apparent when consenting to individual use), and the need to address harms cumulatively and holistically (including how individual choices affect social values more broadly). Automation could make it possible to build tools that can help make the effects of scale and aggregation more transparent and could assist in understanding the broader social implications of practices across different kinds of sectors of activity.

However, we take the view that a focus on consumers and informed consent is too narrow. Privacy policies have a transparency function that goes beyond consumer consent in a number of ways. First, although regulators increasingly stress that informed consent is better operationalized through “layered” policies, visualization techniques, or dynamic permissions²⁵, it remains important to have a comprehensive privacy policy. For example, the current guidelines on the GDPR indicate that “the entirety of the information addressed to data subjects should also be available to them in one single place or one complete document (whether in a digital or paper format) which can be easily accessed by a data subject should they wish to consult the entirety of the information addressed to them”.²⁶ When a privacy policy is no longer the main means of notification to consumers, it is better understood as a mechanism that provides a full account of the authorizations based upon consent rather than as a mechanism for obtaining informed consent.

Second, in most data protection law consent is not the only basis for processing data. There can be other kinds of authorizations and many exceptions to consent. However, there remain transparency obligations in relation to these data flows. Privacy policies are a means for organizations to self-report their data practices in a much more general way than a focus on consent suggests. We can therefore shift away from thinking about the role of privacy policies as enabling informed consent and instead think about these policies as enabling meaningful accountability. This opens up a different way of thinking about work on automating privacy policies.

Third, this shift towards meaningful accountability also requires a shift in the focus on *who* should find privacy policies more usable. Instead of relying upon consumers to hold organizations to account for their privacy practices, either through market choices or through

²⁴ See Pribot, online: https://pribot.org/files/Polisis_Technical_Report.pdf.

²⁵ Office of the Privacy Commissioner of Canada, “Guidelines for obtaining meaningful consent” (May 2018), online: https://www.priv.gc.ca/en/privacy-topics/collecting-personal-information/consent/gl_omc_201805/.

²⁶ Article 29 Data Protection Working Party, *Guidelines on Transparency Under Regulation 2016/679* (2018) 17/EN (WP 260 rev 01) at para. 17, online: https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=622227.

complaints to regulators, we should look to making the role of regulators and other parties more effective. In its 2012 report on consumer privacy, the FTC reported:

While acknowledging privacy policies' current deficiencies, many roundtable participants agreed that the policies still have value – they provide an important accountability function by educating consumer advocates, regulators, the media, and other interested parties about the companies' data practices. Accordingly, Commission staff called on companies to provide clear and concise descriptions of their data collection and use practices.²⁷

There are a range of actors who might be better placed than consumers to hold organizations to account. Automating policies can provide regulators and others parties with tools to understand data practices and compare them at scales that are otherwise difficult to attain.

To give some sense of the problem of scale, consider mobile apps. In December 2017 statista.com reported that the number of apps available for download from Google Play alone was 3.5 M (a doubling since 2015).²⁸ For regulators to simply examine the new apps available from Google Play between December 2016 and December 2017 it would have required reviewing almost 2500 apps every day throughout the year. How can regulators with limited resources scale their capacity to oversee this burgeoning new economy? Regulators can use innovative methods to try to address these problems of sheer scale, such as the Privacy Sweep program undertaken by the Global Privacy Enforcement Network (GPEN). However, we can expect that, despite pooling resources and dedicating blocks of staff time, regulators will continue to fall farther behind.

Shifting from a consent focus to an accountability focus also highlights the importance of integrating privacy policy analysis with methods of auditing organizations for policy compliance. For example, Feigenbaum et al argue that too many privacy and security researchers in technical fields have focused on preventative measures (like the authentication of identity for access to information) and need to now shift to an accountability framework in relation to the online world:

When a policy-governed action occurs, it should be possible to determine (perhaps after the fact) whether an applicable policy has been violated and, if so, to have the violators face appropriate consequences. A move in this direction would make the online world more like the offline world, in which

²⁷ Fed. Trade Comm'n, "Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Business and Policymakers" (March 2012), online (pdf): <https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf>.

²⁸ Jessica Clement, "Google Play: number of available apps 2009-2020" (17 June 17 2020), online: Statista <https://www.statista.com/statistics/266210/number-of-available-applications-in-the-google-play-store/>.

*potential violations of security and privacy are often deterred by the prospect of negative consequences rather than prevented by truly unbreakable locks.*²⁹

To do this requires a way of understanding which actions are policy-governed, what those policies are, and then determining whether there has been a violation. If privacy policies function as a self-reporting mechanism regarding data processing activities and authorities, then they can aid in this shift towards an accountability framework. What is needed in addition are tools to determine compliance with that policy framework.

There is already interesting work on privacy auditing that aims to automate the matching of privacy policies with actual practices. For example, Sen et al have made progress in automating internal privacy audits that ensure that an organization is complying with its privacy policy. The project makes use of “Legalease” — a special language developed for translating privacy policies into encodable policy clauses. The translation into Legalease is done by legal teams and organizational privacy champions. They then use specialized software (Grok) to automate understanding information practices within the organization and match this against the privacy policy permission.³⁰ If we instead use AI tools to automate the reading and classifying of privacy policies, we might be able to skip the step of translating a policy into a special additional language.

In sum, privacy policies are an important aspect of the current infrastructure of data transparency but they suffer from many well-known defects. Many projects to improve the usability of privacy policies focus on improving consumer understanding in order to better facilitate the role of privacy policies in enhancing meaningful consent. In contrast, we argue that we should improve the usability of privacy policies in order to better facilitate their role in enhancing meaningful accountability. We believe that automating policy analysis is a key part of this endeavour and will allow for the creation of additional tools that can help regulators understand data flows at scales otherwise difficult to manage and also provide better means of auditing some practices to see if they in fact conform to the self-reporting provided through privacy policies. However, as the following section outlines, there remain some serious limitations to this strategy despite its promise.

2.2 The Limits of Using Analog Tools to Regulate a Digital World

The benefits of continuing to improve policy automation are that privacy policies are the most widely implemented transparency technology organizations already employ. We find them used across many different areas of activity, and across different jurisdictions. Continuing to find

²⁹Joan Feigenbaum, James A. Hendler, Aaron D. Jagard, Daniel J. Weitzner & Rebecca N. Wright, “Accountability and Deterrence In Online Life (Extended Abstract)” (2011) Yale Univ., online (pdf): <https://dedis.cs.yale.edu/dissent/papers/websci11.pdf>. See also Daniel J. Weitzner, Harold Abelson, Tim Berners-Lee, Joan Feigenbaum, James Hendler & Gerald Jay Sussman, “Information Accountability” (2008) 51 COMM. OF THE ACM at 82, 85.

³⁰Shayak Sen et al., “Bootstrapping Privacy Compliance in Big Data Systems” (2014) Proceedings of the 2014 IEEE Symposium on Security and Privacy 327.

ways to use them to improve the transparency and accountability of data flows in our complex data ecosystem, where data constantly crosses boundaries of organizations, activities, and jurisdictions, is more practical than trying to implement a new system of data transparency. However, there are also important reasons to move beyond a reliance on privacy policies, even if analysis of these policies could be automated. In this section, we outline what these reasons are and then offer several policy proposals for how to improve policy automation as an important pragmatic, medium-term goal.

2.2.1 Moving Beyond Privacy Policies

There are at least three reasons to move beyond a reliance on privacy policies, even if we can succeed in automating their analysis. One reason is that privacy policies are analog tools trying to regulate a digital world. These policies remain essentially a print technology and, as such, their representation of information practices is static. These documents are meant to be what Zimmerman classifies as an *ex ante* transparency tool that provides individuals with knowledge about information practices before data collection and processing.³¹ Part of the difficulty lies in using present language to capture future uses while still ensuring that these future uses will comply with legal requirements. As Sen et al outline,

*Privacy policies are typically crafted by lawyers in a corporate legal team to adhere to all applicable laws and regulations worldwide. Due to the rapid change in product features and internal processes, these policies are necessarily specified using high-level policy concepts that may not cleanly map to the products that are expected to comply with them.*³²

This generates privacy policies that are long and complex in order to comply with legal requirements but which still do not necessarily provide enough specificity to understand specific information practices. However, complex information ecosystems are dynamic, rather than static. Because of this, privacy policies will always have problems providing transparency regarding future uses of data.

Privacy policies also do not provide a clear picture of privacy “defaults”. For example, Facebook’s Data Policy states: “When you share and communicate using our Services, you choose the audience who can see what you share.”³³ This indicates what is in the user’s choice in relation to data sharing but it does not help the user -- or anyone else -- to analyze the initial default settings. Facebook sets a default for different types of information, such as “friends” and then allows individuals to change these settings and choose a different audience. However, defaults matter and can nudge an individual to make a privacy choice that is not consistent with his or her privacy preferences or that raises issues of broader social concern.

For these reasons, rooted in the static nature of privacy policies, even automated privacy policies cannot alone “solve” the data transparency problem. We might be better off trying to

³¹ Christian Zimmerman, “A Categorization of Transparency Enhancing Technologies” (2015), online: <https://arxiv.org/abs/1507.04914>, at 7.

³² Shayak Sen et al., *supra* note 30.

³³ See Facebook, “Data Policy”, online: Facebook <https://www.facebook.com/policy.php>.

think of alternative means of creating what is essentially a repository of information, but one that could be made more dynamic. As a contrast to a policy, consider a database. A database holds data and allows its data to be accessed and represented in different ways for different purposes and users. By automating the reading and classifying of existing privacy policies we can create some of this functionality. However, a further shift that could happen is a shift away from using humans to manually “map” a complex information ecosystem and then translate this into the static language of a privacy policy. Instead, this “mapping” should be automated through technological means with different user tools for transparency then built on top of this automated mapping. Instead of having different forms of transparency that serve different users -- consumers, regulators, developers -- there would be one underlying form of transparency (the database), this form could be audited to verify that it indeed does map onto actual information practices, and then different users can have different tools built for accessing this information that is tailored to their needs. This is important, because different users have different needs. As Brennan-Marquez and Susser argue, providing greater “under-the-hood” information that is important for compliance purposes can confuse individual users as much as it can enlighten.³⁴ Individuals, regulators, and developers all have different informational needs.

A second reason to move past reliance on privacy policies for transparency disclosures is that complex exchange networks will likely require better technical means to track data flows and these will open possibilities to create new forms of transparency and accountability. One effort that moves in this direction is the IAB Europe Transparency and Consent Framework (TCF).³⁵ The extent to which the TCF addresses the consent and transparency issues associated with the online ad industry is debatable and currently part of a broad GDPR challenge.³⁶ What it does is provide consumers with the ability to view which organizations might get access to their data within the advertising ecosystem, and provide consent based upon more fine-grained choices relating to purposes and organizations. What it does not do is provide technical means to audit recipient organizations for unauthorized data uses.³⁷ Managing this requires developing technologies to manage complex information flow constraints that go beyond current methods of binary access controls.

Finally, a third reason for moving beyond privacy policies is that part of the problem with accountability is the number of organizations that have data and the complexity of purposes for which they use that data. One way to address this problem is to stop releasing data to multiple third parties. We have already outlined how difficult it is to manage residual risks of misuse through contracts like data sharing agreements. A different way to handle complex data flows is to provide controlled access to data within a trusted computing environment where uses are

³⁴ Kiel Brennan-Marques & Daniel Susser, “Obstacles to Transparency in Privacy Engineering” (2016) IEEE Security and Privacy Workshops 49.

³⁵ See generally Jennifer Derke, “Transparency & Consent Framework” (24 April 2019), online: *Github* <https://github.com/InteractiveAdvertisingBureau/GDPR-Transparency-and-Consent-Framework/blob/master/pubvendors.json%20v1.0%20Draft%20for%20Public%20Comment.md#goals>.

³⁶ Johnny Ryan, “Formal GDPR complaint against Google’s internal data free-for-all” (16 March 2020), online: *Brave* <https://brave.com/google-internal-data-free-for-all/>.

³⁷ Johnny Ryan, “Google and IAB’s inadequate proposals to reform RTB” (21 Jan 2020), online: *Brave* <https://brave.com/google-iab-reform/>.

logged and auditable rather than disclosing data to the third parties who want to use it.³⁸ This allows for the development of different methods to ensure transparency and policy compliance within that trusted computing environment.

However, these broader shifts away from using privacy policies as a mechanism for data transparency will take time and experimentation to develop. Although the automation of privacy policies is not a magic bullet, it has a place in enhancing meaningful accountability in the medium-term. As already mentioned, one of its main advantages is that most jurisdictions already require organizations to have a privacy policy and so automating privacy policy analysis makes use of a significant existing infrastructure and seeks to make important improvements.

2.2.2 Improving Policy Automation

Despite the practical advantages of utilizing an existing infrastructure of privacy policies, properly realizing the medium-term goals of policy automation will require further technical research. However, even as a medium-term solution, an acceptable level of policy automation cannot be realized through technical means alone. We think that a number of potential policy reforms can also help to improve the transparency function of privacy policies and their better integration with audit mechanisms. We outline several potential reforms here: standardization, an emphasis that the goal of standardization is to enhance automation rather than immediate consumer comprehension, adequate audit powers for regulators, and ensuring that such tools conform to norms of fairness.

Some of the limits of automation can be addressed through ongoing research to improve the automation of policy analysis, such as through work that seeks to understand and address the impact of linguistic ambiguity on automation.³⁹ However, if privacy policies are poorly drafted and inherently vague then automating analysis is no solution. The first reform we propose is therefore legal requirements to standardize elements of policies. The GDPR prescribes some information that must be provided to the data subject but other jurisdictions rely upon guidelines.⁴⁰ Even the GDPR rules are fairly general and leave open the possibility of using the kind of vague language that makes aspects of interpreting privacy policies difficult. For example, the Article 29 Data Protection Working Party guidance on transparency outlines the need for clear and plain language but leaves open how to take up this advice.⁴¹ The FTC has called on industry to create more standardized policies, leaving this to self-regulation.⁴²

³⁸ See, eg, Lisa Austin & David Lie, “Safe Sharing Sites” (2019) 94 NYU L Rev 581.

³⁹ See, eg, Jaspreet Bhatia, Travis D. Breaux, Joel R. Reidenberg, and Thomas B. Norton, “A Theory of Vagueness and Privacy Risk Perception” (2016) IEEE 24th International Conference on Requirements Engineering 26; Sebastian Zimmeck & Steven M. Bellovin, *supra* note 20; Qiao Zhang, “Fuzziness-vagueness-generality-ambiguity” (1998) 29 J. of Pragmatics, at 13.

⁴⁰ *GDPR*, *supra* note 9, arts. 13-14; Office of the Privacy Comm’r of Canada, “Guidelines For Obtaining Meaningful Consent” (May 2018), online: https://www.priv.gc.ca/en/privacy-topics/collecting-personal-information/consent/gl_omc_201805/ [hereinafter “Consent Guidelines”].

⁴¹ Article 29 Data Protection Working Party, *supra* note 26.

⁴² Fed. Trade Comm’n, *supra* note 27

A model for standardization is the Model Privacy Form - which was released to aid companies in complying with the Gramm Leach Bliley Act (GLBA).⁴³ A study by Cranor looked at US financial institutions' policies that conform to this Model Privacy Form.⁴⁴ This study used a simple automated parsing algorithm to extract information from the policies and managed to reach 100% accuracy on nearly all sections.⁴⁵ Reidenberg et al also did a study looking at various privacy policies from the perspective of a vagueness taxonomy in order to determine the impact of different regulatory regimes.⁴⁶ The study looked at policies written to be compliant with three different regulatory schemes: no regulation, EU-US Safe Harbour regulations, and financial policies written using the Gramm Leach Bliley Model Privacy Form. The results showed that unregulated policies had the most vagueness and those written in line with the Model Privacy Form had the least amount of vagueness.

We think there are benefits to standardization but that we need to shift the framing of these benefits away from a focus on enhanced consumer understanding to an enhanced ability to automate policy analysis. Therefore, the second reform we propose is to ensure that the standardization of data disclosures supports automation. Once we can better automate this analysis then we can create many different kinds of tools that can enhance consumer understanding as well as provide investigative tools for regulators. In speaking about the US Security Exchange Commission's use of machine learning tools to detect misconduct, Bauguess noted that

*[t]he success of today's new technology [for detecting misconduct] depends on **the machine readability of decision-relevant information**. And I don't mean just for numerical data, but for all types of information. This includes narrative disclosures and analyses found in the written word. It also includes contextual information about the information, or data about the data, often referred to as "metadata." Today's advanced machine learning methods are able to draw incredibly valuable insights from these types of information, but only when it is made available in formats that allow for large-scale ingestion in a timely and efficient manner.*⁴⁷

The US SEC has proposed a form of filing that combines a human readable format with and eXtensible Business Reporting Language (XBRL) format. While financial disclosures are

⁴³ 12 C.F.R. § 1016 (2018) (appendix).

⁴⁴ Lorrie Faith Cranor, Pedro Giovanni Leon & Blase Ur, "A Large-scale Evaluation of US Financial Institutions' Standardized Privacy Notices" (2016) 10 Acm Transactions on the Web; Peter Swire, "The Surprising Virtues of the New Financial Privacy Law" (2002) 86 Minn L Rev 1263, 1315-16.

⁴⁵ The algorithm was a simple parsing algorithm, which is likely much less computationally costly than the ML algorithms that were used to analyze the unstandardized notices – with much less success.

⁴⁶ Joel R. Reidenberg et al., "Ambiguity in Privacy Policies and the Impact of Regulation" (2016) 45 J Legal Stud S163.

⁴⁷ Scott W. Bauguess, Deputy Chief Economist and Deputy Director, Division of Economic and Risk Analysis, "Keynote Address at the Financial Information Management (FIMA) Conference 2018: The Role of Machine Readability in an AI World" (May 3, 2018), online: *US Securities and Exchange Commission* <https://www.sec.gov/news/speech/speech-bauguess-050318>.

different from the data disclosures we are contemplating, privacy regulators should look to these examples in other areas.

Transparency is not enough to ensure meaningful accountability. This is particularly true when the transparency is accomplished through self-reporting, as it is with privacy policies. Privacy law needs to adopt lessons from other regulatory contexts that rely on forms of self-reporting. For example, tax law relies upon the self-reporting of individuals and businesses but then has methods of auditing.⁴⁸ We think that we should shift away from a focus on whether consumers can understand privacy policies and either vote with their feet by rejecting services with privacy-invasive practices or enforcing their rights through bringing a complaint to a regulator and instead focus on whether regulators can understand privacy policies and audit for compliance. Therefore, the third needed reform to support automation is to ensure that regulators have adequate audit powers. For example, maybe regulators should be empowered to make random audits of companies. The GDPR introduced a principle of demonstrable accountability and this could be a further enhancement of the idea. Being able to demonstrate accountability could involve being able to meet the conditions of a data audit. Some of this already occurs on a voluntary basis, as when organizations choose to use something like the TRUSTe certification.⁴⁹ In the financial sectors we developed methods of accounting that facilitate various forms of financial auditing, such as the US Generally Accepted Accounting Principles (GAAP) and the International Financial Reporting Standards (IFRS). We need to stop placing the emphasis on the role of the consumer and see that the data disclosures in a privacy policy are as important to accountability in the data economy as financial disclosures are. This means regulating these disclosures, imposing standards, and matching this with better auditing.⁵⁰

In addition to these three policy reforms that support the automation of privacy policy analysis and auditability of data flows, more thinking needs to be done to ensure that the use of such tools by regulators is consistent with emerging norms regarding fairness questions pertaining to the use of machine learning technologies in administrative processes more generally.⁵¹ What is demanded in any particular context will depend upon the nature of the

⁴⁸ For a review of the literature on tax compliance, concluding that auditing coupled with sanctions is generally an effective deterrent, see Leandra Lederman, “Does Enforcement Reduce Voluntary Tax Compliance?” (2018) 2018 BYU L Rev 623.

⁴⁹ See “TRUSTe Data Collection Certification” online: *TrustArc* <https://trustarc.com/truste-certifications/data-certification/>.

⁵⁰ For related proposals for auditing in relation to the use of algorithms by business and government, see Danielle Keats Citron & Frank Pasquale, “The Scored Society” (2014) 89 Wash L Rev 1; Deven R. Desai & Joshua Kroll, “Trust by Verify: A Guide to Algorithms and the Law” (2017) 31 Harv J L & Tech 1; Kate Crawford & Jason Shultz, “Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms” (2014) 55 BC L Rev 93.

⁵¹ See generally Ryan Calo & Danielle Keats Citron, “The Automated Administrative State: A Crisis of Legitimacy” (2020) 70 Emory Law J 797; Deirdre K. Mulligan & Kenneth A. Bamberger, “Procurement as Policy: Administrative Process for Machine Learning” (2019) 34 Berkeley Tech L J 773; Solon Barocas & Andrew D. Selbst, “Big Data’s Disparate Impact” (2016) 104 Calif. L. Rev. 671; Citron & Pasquale, *supra* note 50; Joshua A. Kroll et al., “Accountable Algorithms” (2017) 165 U Pa L Rev 633; Crawford & Shultz, *supra* note 50; Danielle Keats Citron, “Technological Due Process” (2008) 85 Wash U

proposed tool and here we only offer a few general observations. First, automated tools for regulators -- like the prototype we created in our AppTrans project discussed in the following section -- is not meant to automate findings of non-compliance but to function as a flag for regulators indicating areas where they can focus their energies in relation to investigations or education. This still raises issues about the need to assess such tools for accuracy, to make these assessments open to public scrutiny, and to train users so that they can understand the limits of the tools for their purposes. Making such tools open source can also be a way of enabling public review of these tools and enabling public trust in their use.⁵² Second, automated tools could help with more targeted investigations but, again, not by replacing human decision-makers but by providing them with important information. Information asymmetries like those in the technology sector can hinder the work of regulators and provide a justification for forms of regulatory monitoring; this monitoring need not be punitive but could be part of collaborative governance models.⁵³

3 Automating Privacy Policy Analysis

3.1 The AppTrans Project

In this section we describe a project that we undertook to demonstrate the possibility of combining policy automation with the verification of data flows. Our project -- AppTrans -- created a prototype of a tool that could be used by regulators to determine when a mobile application was collecting personal information that it was not declaring in its privacy policy.⁵⁴ We outline our methodology, its limitations, and what we learned for future iterations. After creating the AppTrans tool, we used it to analyze over 700 mobile apps and their accompanying privacy policies, we found that approximately 60% of mobile apps tested were collecting personal information that was not declared in their privacy policies. The following section discusses these findings in more detail.

The creation of the AppTrans tool involved three main steps. The first was to automate privacy policy analysis. The second was to automate the analysis of data flows -- a kind of auditing function that would allow us to understand whether a mobile app actually collects personal information as distinct from what it says its practices are. The third was to automatically compare the privacy policy to the data flows in order to determine whether there were data flows that were not declared in the privacy policy -- allowing for verification of the claims made in the privacy policy. The basic idea was to create a tool that could automate this

L Rev 1249; Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Cambridge, Mass.: Harvard University Press, 2015).

⁵² Citron, "Technological Due Process", *supra* note 51; The AppTrans source code is available here: <http://www.eecg.toronto.edu/~lie/downloads/opc-2018.tar.gz>.

⁵³ See, eg, Rory Van Loo, "The Missing Regulatory State: Monitoring Business in an Age of Surveillance" (2019) 72 Vand L Rev 1563, at 1580, 1621.

⁵⁴ See also Sebastian Zimmeck et al., "Mobile App Privacy Compliance: Automated Technology to Help Regulators, App Stores and Developers" (2017) Proceedings of the Thirteenth Symposium on Usable Privacy and Security.

kind of analysis and perform it at scale so that regulators could scan large groups of applications and flag potentially problematic practices for follow-up. We developed AppTrans as a research prototype and it was not intended for broad deployment and use. Nonetheless, our results show that this type of tool is feasible and can offer regulators new insights into the activities that they regulate.

Our method for automating privacy policy analysis was to use supervised machine learning to create a model. Supervised machine learning is a type of machine learning that uses a labeled dataset as “training data”. For example, by being shown many different examples of text with particular labels, we can train a model that will then be able to predict the appropriate labels when shown unlabeled text.

For our project, using supervised machine learning meant that we had to create a database of privacy policies and then use humans to label segments of these policies. We constructed a training set of 32, 808 privacy policies by “crawling” (using a software program to automatically scan and extract the relevant information) the Google Play store and then cleaned up this dataset by removing invalid documents, documents that were not actually privacy policies, and non-English privacy policies. We then used a paid crowdsourcing service -- Amazon’s Mechanical Turk service (MTurk) -- to label the policies. Amazon MTurk workers were asked to label segments of text that contained information about data collection by what type of data was collected. We focused on 3 classes of information: location, contact information, and device identifiers. AppTrans initially focused on only these three classes of information as they are the most common, but AppTrans can be extended to other classes of information in a straightforward manner.

To ensure accuracy of labelling, we developed a protocol for testing the MTurk labelling.⁵⁵ We used law students to label some policy segments and then included some of these “test” policy segments in the policy segments that the MTurk workers were asked to label. If the MTurk workers labelled these test policies incorrectly, then we excluded their results. Further, each text segment was given to 5 different workers to label and we only used a label if 4/5 or more of the workers gave the same label. This methodology resulted in a labeled dataset of 2,254 policy segments about data collection. We then trained a machine learning model on this dataset. This model allows us to provide it with an unlabelled privacy policy and automatically tell us whether that privacy policy declares a collection of any of 3 classes of information. For further quality control, we evaluated our model against an existing manually labeled set.⁵⁶ Our findings were that our model can correctly classify whether a privacy policy declares collection of one of the 3 information classes with 95% accuracy.⁵⁷

The second component of our tool involved determining what personal information is collected by the application. We did this by analyzing the application’s software code to

⁵⁵ Since this component involved human subjects, this component of the research was approved by the University of Toronto’s research ethics board.

⁵⁶ Shomir Wilson et al., “The Creation and Analysis of a Website Privacy Policy Corpus” (2016) Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics 1330.

⁵⁷ For a different approaches to labeling privacy policies and their accuracy, see Florencia Marotta-Wurgler, “Understanding Privacy Policies: Content, Self-Regulation, & Markets” (2016) 45 *Journal of Legal Studies* S13 and M Lippi et al, “Claudette: an automated detector of potentially unfair clauses in online terms of service” (2019) 27 *Artificial Intelligence and Law* 117.

determine which mobile phone “permissions” the application was able to access and whether the information that it collected through such permission was sent off the phone to a remote server. Mobile phone permissions are the way in which applications can access information through interacting with the phone’s operating system and therefore other basic features of the phone. For example, apps might access a mobile phone’s contacts, camera, location, microphone, or many other types of permissions. From the type of permission that the app was seeking to access, we could infer that it was collecting personal information of a particular type. We considered this a collection when the information was sent off the phone in order to eliminate the case where the users could store information in the application but the app company itself (or other third parties) did not have access to it.

It is important to note that information collected through permissions is different than information that a user might input through the app’s user interface. There is currently no good method for automatically determining whether the inputs to a user interface are personal information and whether this information is sent off-device. When we did a manual review of 19 policies that AppTrans had flagged as non-compliant, we found one instance where the policy declared that it collected location data but our data flow analysis indicated that the app did not collect this data. We believe this discrepancy was due to the information being collected through the user interface, which AppTrans cannot detect.

We automated our data flow analysis using an existing tool – the FlowDroid static analysis tool (version 25)⁵⁸. Static code analysis is a method of analyzing software that does not involve running the software to see how it performs with users in real time. We note that no tool, including FlowDroid, can detect all possible flows that may exist.⁵⁹ In a study performed by the FlowDroid authors on their own tool, they found that FlowDroid could detect 93% of flows correctly while having a false reporting rate of 14%. Other research projects on mobile app privacy have successfully used dynamic, rather than static, analysis, which involves analyzing software while it is being used. where software is analyzed while it is being used.⁶⁰ Static analysis is faster and so is able to analyze more applications in more detail with the same amount of time and compute resources, but may give less precise results. We therefore felt that static analysis is more amenable to the broad study of mobile applications and privacy policies we were aiming for here. In particular, compared to a previous study that used dynamic analysis

⁵⁸ Steven Arzt et al., “FlowDroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps” (2014) Proceedings of the 35th ACM SIGPLAN conference on programming language design and implementation 259.

⁵⁹ To understand this, an analogy can be made to the well-known halting problem, which says that it is impossible to write one program that will in all cases be able to determine whether another program halts or runs forever. Similarly, the same result can be used to show that one cannot determine with certainty whether a program will perform any particular action. See Alonzo Church, “An Unsolvability Problem of Elementary Number Theory” (1936) 58 Am J Math 345.

⁶⁰ See, eg, Joel Reardon et al., “50 Ways to Leak Your Data: An Exploration of Apps’ Circumvention of the Android Permissions System” (2019) Proceedings of the 28th USENIX Security Symposium; Serge Egelman et al., “Android Permissions Remystified: A Field Study on Contextual Integrity” (2015) Proceedings of the 24th USENIX Security Symposium.

to detect violations of the *Children's Online Privacy Protection Act* (COPPA) in the US⁶¹, we were able to give information about specific flows (i.e. release of contact information), while the COPPA study used dynamic analysis just to cause the app to generate network traffic, which was subsequently analyzed for the presence of personal information. In general, dynamic analysis and static analysis have complementary strengths and weaknesses and it is possible to combine the two to achieve both more comprehensive and precise results at the cost of more computing resources and time. Therefore it is possible that some of these methods could be incorporated into other future iterations of this kind of tool.

As it was intended as an exploratory prototype, the AppTrans project was also limited to certain aspects of privacy policies and mobile applications. For one thing, AppTrans focused exclusively on the Android platform and on Android Mobile Applications. This was for three reasons: 1) Android dominates the smartphone market globally and there are many applications; 2) the existence of the Google Play store constitutes a central place where application code and their associated privacy policies can be easily attained; and 3) open-source tools are readily available for the Android platform making prototyping easier. Another limitation was that AppTrans focused on personal information collection as opposed to use or disclosure. This was for two main reasons: 1) collection is often more clearly and unambiguously declared in privacy policies than use or disclosure; and 2) collection must be done by the mobile application itself whereas use or disclosure can involve communication with a backend server and AppTrans could not access the source code of the backend server.

In the following section we outline our findings when we used the AppTrans tool to analyze over 757 mobile applications and their privacy policies. However, quite apart from what we learned by deploying our tool, we can make several important conclusions about the general feasibility of creating such tools. Our initial experience developing the AppTrans prototype suggests that automated code analysis and machine learning together can form a basis for building useful and reliable tools for automating some aspects of detecting compliance problems. However, as we will discuss in relation to our follow-on work, machine learning relies on having clear and unambiguous labels on text, and obtaining these at scale was the major challenge in the AppTrans work. In addition to the problem of scale is the problem of vagueness and ambiguity in the policies themselves. As we discuss in relation to our follow-on work, although privacy policies are relatively clear regarding data collection they lack clarity regarding other aspects of data flows. Another challenge is obtaining access to the code of the software handling the personal information to check for non-compliance. In AppTrans, we could only obtain the code of the mobile applications, as they must be downloaded and installed on a mobile phone to operate, but often the personal information is transferred to a service running on the Internet, for which the code is not available. Further, service code can be more diverse and difficult to analyze than mobile application code, which can make the code analysis problem significantly more challenging even if access is available. These are all reasons why automated tools like AppTrans can improve upon our current transparency environment but at the same

⁶¹ Irwin Reyes et al., "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale" (2018) 3 Proc Privacy Enhancing Techs 63.

time point to why we need to move towards more self-consciously engineering our digital environment to enable more robust forms of transparency and accountability.

3.2 Our Findings and The Third Party Problem

3.2.1 AppTrans Findings

We used our AppTrans tool to evaluate a large number of mobile applications in order to determine whether any of these apps were collecting personal information (location, contacts, device ID) that they were not declaring in their privacy policy. Of the more than 700 apps and policies that we analyzed, we found that approximately 60% of tested apps were likely in violation of legal transparency requirements. In order to drill further into the problem, we determined that it was third party code that was responsible for 85% of the problematic applications. This means that while the code written by the application developer complies with the app’s privacy policy, in the majority of cases non-compliance occurs because there is collection of data by third party code that is not declared in the privacy policy. Third party code is software code created by third parties such as analytics or advertising companies and is incorporated by application developers in order to make use of these services.

The details of our findings are as follows:

	Location	Contacts	Device ID	Average
Policy-app pairs	404	19	334	
Non-compliance	217 (53.7%)	12 (63.2%)	206 (61.7%)	59.5%
1st party	26 (12.0%)	2 (16.7%)	24 (11.7%)	13.4%
3rd party	186 (85.7%)	10 (83.3%)	179 (86.9%)	85.3%
1st and 3rd	5 (2.3%)	0 (0.0%)	3 (1.5%)	1.3%

We acquired our test set of mobile applications and privacy policies from the Playdrone dataset, which is a large index of Android applications including their source code.⁶² Since the

⁶² Nicolas Viennot, Edward Garcia & Jason Nieh, “A measurement study of Google Play” (2014) 42 ACM Sigmetrics Performance Evaluation Rev. 221.

Playdrone dataset is from 2014 we used Archive.org to retrieve old snapshots of privacy policies so that we could ensure that we would only use application-privacy policy pairs that differ by less than 60 days in age. We then filtered these pairs to only include applications that sought access to location, contacts or device identifiers (as determined through analyzing their software code). This resulted in a dataset of 757 privacy policy-application pairs.

3.2.2 The Third Party Problem

The third party result is striking. A deeper understanding of its implications requires interrogating further the role of third party code in mobile applications.

From the perspective of the user, a mobile application appears to only involve a relationship between the application and the user interacting with it. However, the incentives and methods of monetization in the mobile application ecosystem complicate this seemingly simple picture. This is because monetization in the application ecosystem either involves a) advertising being placed in the mobile application by the application developer and viewed by the user or b) in-app payments by the user for additions to or features in the mobile application after installation time. Both of these methods of monetization encourage application developers to involve third parties, in the form of third party code libraries, in their applications. In the case of advertising the relationship is fairly simple: the third party advertising library collects information about the user and serves “targeted advertising” to the end user. The situation with in-app payments is slightly more complex as the amount of in-app payments is related to the amount of time the user spends using the application---a property often called “engagement” in the mobile application industry.⁶³ To maximize engagement, application developers benefit from having detailed demographic information about their users, as well as behavioral information about how, when and where users use their applications. Because the collection of this information is complex and requires data collection and storage infrastructure, application developers often use third party analytics libraries to collect, store and analyze user information to help guide their application development.

In both these cases, the third party advertising and analytics libraries are incorporated into the application in such a way that they are effectively part of the application, and have all the capabilities of the application. On a mobile smartphone, the operating system (Android and iOS) is responsible for determining whether an application may access sensitive information, such as the user’s location, browsing history, contact list, etc. Normally, this access is granted or denied under the direction of the smartphone user. However, because the third party library is incorporated directly into the application, current smartphone architectures do not allow a user to indicate to whom they are granting a particular permission (for example, the ability to access their current location) -- to the application, the library or both. It is important to note that as a consequence of the third party code being implemented as a library, there is no way to separate third party access to permissions from the app developer’s access. Whatever permissions the app has, the library has. If the user wishes to disable a permission so that a third party library

⁶³ App Annie, “The State of Mobile 2019” 27 online: *AppAnnie* <https://www.appannie.com/en/go/state-of-mobile-2019/>; See, eg Helen Vakhnenko, “How to Increase Mobile App Engagement - Complete Guide” online (blog): *Agilie* <https://agilie.com/en/blog/how-to-increase-mobile-app-engagement-complete-guide>.

cannot use it, the user has to disable the permission for the entire app and will lose whatever app functionality is associated with the permission. If an application wants to constrain third party access then it must either tamper with the library – which would almost certainly violate the library’s terms of service – or do so through non-technical means like contracts, an issue we discuss further in s. 4.

The use of third parties to analyze user behaviour and to deliver advertising is a ubiquitous, opaque, and confusing aspect of the practices of both websites and apps. Initiatives like Vermont’s privacy law have attempted to shine light on the shadowy world of data brokers while legal complaints under the GDPR are currently attempting to address the online advertising industry.⁶⁴ Although many people have a passing familiarity with web tracking through the use of cookies and similar technologies, most people have less understanding of the role of third parties in the mobile app universe.⁶⁵ Recently, the widespread use of third party code to track users has received some high-profile media attention. The New York Times and the Washington Post have both reported on the issue, how prevalent it is, and how consumers are largely in the dark regarding the resulting data flows and their implications.⁶⁶

Third party libraries are often complex and non-transparent to developers as well, so developers might incorporate third party code without fully understanding what this means for data flows. To make matters worse, the relationship between application developers and third

⁶⁴ Vt. Stat. Ann. tit. 9, § 2445 (2019); for background on the law, see generally Brittany A. Martin, “The Unregulated Underground Market for your Data: Providing Adequate Protections for Consumer Privacy in the Modern Era” (2020) 105 Iowa L. Rev. 865; Steven Melendez, “A landmark Vermont law nudges over 120 data brokers out of the shadows” *Fast Company* (2 March 2019), online: <https://www.fastcompany.com/90302036/over-120-data-brokers-inch-out-of-the-shadows-under-landmark-vermont-law>; Fed. Trade Comm’n, “Data Brokers: A Call for Transparency and Accountability” (May 2014), online (pdf): <https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf>.

Regarding GDPR complaints, see, eg, Johnny Ryan, “Regulatory complaint concerning massive, web-wide data breach by Google and other “ad tech” companies under Europe’s GDPR” (12 September 2018), online: *Brave* <https://brave.com/adtech-data-breach-complaint/>; See generally Natasha Lomas, “Google and IAB ad category lists show ‘massive leakage of highly intimate data,’ GDPR complaint claims” *Techcrunch* (28 January 2019), online: <https://techcrunch.com/2019/01/27/google-and-iab-ad-category-lists-show-massive-leakage-of-highly-intimate-data-gdpr-complaint-claims/>.

⁶⁵ Chris Jay Hoofnagle et al., “How Different are Young Adults from Older Adults When it Comes to Information Privacy Attitudes and Policies?” (2010) [unpublished], online: *SSRN* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1589864.

⁶⁶ Geoffrey A. Fowler, “In the Middle of the Night, your iPhone is Busy Sharing your Personal Data” *Washington Post* (29 May 2019) A14; Sam Schechner & Mark Secanda, “You Give Apps Sensitive Personal Information. Then They Tell Facebook” *Wall Street Journal* (22 February 2019), online: <https://www.wsj.com/articles/you-give-apps-sensitive-personal-information-then-they-tell-facebook-11550851636>; Jennifer Valentino-DeVries, Natasha Singer, Michael H. Kellerr, and Aaron Krolik, “Your Apps Know Where You Were Last Night, and They’re Not Keeping It Secret” *New York Times* (10 December 2018), online: <https://www.nytimes.com/interactive/2018/12/10/business/location-data-privacy-apps.html>.

parties does not encourage transparency. Since application developers are paid for showing advertisements from the third party, third parties must protect themselves from malicious application developers who may seek to subvert the third party code in order to artificially inflate the number of advertisements shown for monetary gain. As a result, third parties often seek to further obfuscate or obscure the behavior of the third party code.

3.2.3 Transparency and Consent Obligations Regarding Third Parties

We found some evidence that application developers are not necessarily fully aware of their responsibilities for transparency with respect to third party code. For example, we found instances of privacy policies that declared that a mobile application did not collect a category of personal information but that third parties might; in some cases, they provided a link the privacy policy of the third party. One hypothesis, therefore, for why we found such a low level of disclosure regarding third party data collections in the AppTrans project is that the app developers believed that they were not the ones collecting data and therefore did not have transparency obligations. As an example, consider the following segment from one of the policies:

Some content or applications, including advertisements, on the Website are served by third-parties, including advertisers, ad networks and servers, content providers and application providers. These third parties may use cookies alone or in conjunction with web beacons or other tracking technologies to collect information about you when you use our website. The information they collect may be associated with your personal information or they may collect information, including personal information, about your online activities over time and across different websites and other online services. They may use this information to provide you with interest-based (behavioral) advertising or other targeted content. We do not control these third parties' tracking technologies or how they may be used. If you have any questions about an advertisement or other targeted content, you should contact the responsible provider directly.

Aside from the segment being about a website rather than a mobile app, the developer is informing the user about possible collection, but appears to indicate that it is the responsibility of the user to deal with the third party directly to find out what information is being collected and sent to the third party.

In contrast, we argue that app developers have an obligation to disclose third party collection within their own privacy policy. A mobile application's use of a third party library for advertising or analytics is analogous to a website's use of cookies for these purposes. A cookie is a piece of information from a web server that is stored by a user's web browser. Once sent to the user's browser, the cookie is sent to the web server on all subsequent requests. Thus, a cookie can be used to track a user's activity across requests. When combined with a "web beacon", where an element (like an image for example) from one web server (i.e. a tracker) is included on another website then, the tracker can track a user's activities across the websites that have included the tracker's web elements. Cookies have been a controversial element of

online activities, especially in the context of online behavioral advertising because their use has enabled the tracking, profiling, and targeting of individuals.⁶⁷ However, regulators have taken the position that websites need to disclose the use of cookies in their privacy policies, including third party cookies.

Regulators in Europe, Canada, and the US have all taken a broad approach to the question of whether cookies are personally identifiable information (PII) and as such, cookies generally fall under the relevant consumer privacy laws.⁶⁸ The US and Canadian approach is similar and has been to look at the role of the data collection in profiling individual consumers rather than analyzing the nature of the information in abstraction from its use. In its self-regulation guidelines for online behavioural advertising, the FTC disputed the ongoing relevance of the distinction between PII and non-PII and included “any data collected for online behavioral advertising that reasonably could be associated with a particular consumer or with a particular computer or device”.⁶⁹ For similar reasons, the OPC has considered the data used in online behavioural advertising to be “personal information” for the purposes of Canadian private sector data protection law.⁷⁰ Europe takes a stricter view. For example, in its recent decision regarding cookies, the EU Court of Justice invoked the EU ePrivacy Directive in response to the issue of whether cookies were “personal data”:

according to which any information stored in the terminal equipment of users of electronic communications networks are part of the private sphere of the users requiring protection under the European Convention for the Protection of Human Rights and Fundamental Freedoms. That protection applies to any information stored in such terminal equipment, regardless of whether or not it is personal data, and is intended, in particular, as is clear from that recital, to protect users from the risk that hidden identifiers and other similar devices enter those users' terminal equipment without their knowledge.”⁷¹

⁶⁷ Laura J Bowman, “Pulling Back the Curtain: Online Consumer Tracking” (2012) 7 I/S: J. L. & Pol’y For Info, Soc’y 721; Chris Hoofnagle et al., “Behavioral Advertising: The Offer you Cannot Refuse” (2012) Harv L & Pol’y Rev 273.

⁶⁸ The US uses the term “personally identifiable information”, Canadian law uses the term “personal information” and EU law refers to “personal data”.

⁶⁹ Fed. Trade Comm’n, “Self-Regulatory Principles for Online Behavioral Advertising” (February 2009) at 25, online (pdf): <https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-staff-report-self-regulatory-principles-online-behavioral-advertising/p085400behavadreport.pdf>.

⁷⁰ Office of the Privacy Comm’r of Canada, “Policy Position on Online Behavioural Advertising” (2015), online: https://www.priv.gc.ca/en/privacy-topics/advertising-and-marketing/behaviouraltargeted-advertising/bg_ba_1206/.

⁷¹ Case C-673/17, Bundesverband der Verbraucherzentralen und Verbraucherverbände — Verbraucherzentrale Bundesverband eV v. Planet49 GmbH, ECLI:EU:C:2019:801, ¶170 (Oct. 1, 2019), online: <http://curia.europa.eu/juris/document/document.jsf?text=&docid=218462&pageIndex=0&doclang=EN&mode=lst&dir=&occ=first&part=1&cid=9439569>.

All three jurisdictions, despite the differences in their legal frameworks, consider the use of cookies to raise privacy concerns and they regulate it.

Even more importantly from the perspective of our AppTrans project, these different jurisdictions all require websites to disclose cookie use in their privacy policies. Sometimes this is discussed in relation to questions of choice or consent.⁷² However, even where the issue is the use of third party cookies and the third party is required to obtain consent, the first party (the website) must still provide users with notification.⁷³ And in fact Canadian and European rules regarding transparency are getting stricter in relation to generally disclosing the identities of parties who might be recipients of personal information. The GDPR requires the disclosure of the “recipients or categories of recipients”.⁷⁴ The Office of the Privacy Commissioner of Canada has indicated that for the purposes of obtaining meaningful consent, a key element of notification is information about the parties with which personal information is being shared.⁷⁵ In light of the analogies between the use of cookies on websites and the practice of third party data sharing for analytics and advertising in mobile applications, it is reasonable to argue that applications need to disclose any third party data sharing in its privacy policy.

In addition to the obligations of app developers, there is a question of the obligations of the third parties, independently considered. For example, if a mobile application adopts the Google Mobile Ads SDK in order to use AdMob to deliver ads, then Google also needs to ensure that it has the proper authorization for this use, and also that it discloses this information to its advertising partners in compliance with data protection law requirements. We take up this issue in our discussion of the Facebook investigations in section 4, below.

Even if transparency were increased, the role of third parties puts considerable strain on legal and theoretical models of privacy that emphasize individual consent and control. Decisions about which third parties to use and for what purposes are made by the primary organization (e.g. the application) and not the consumer. This does not change with greater transparency -- a mobile application could potentially offer a consumer a choice between receiving targeted advertising or not (as Apple now does) but it is impractical to offer individuals discrete choices regarding the identity of the third party who controls that advertising because the application must make those decisions for multiple users. Moreover, the individual user would still have difficulty understanding the systematic effects of these different choices. For example, a consumer might have three apps on her phone and all might use the same advertising library. Even if data collection for targeted advertising was disclosed to the consumer, and she accepted this for each individual app, she would not necessarily know that all three use the same advertising library. And it would be this third party who would then be able to track her activity across three different apps and potentially profile her in order to deliver those ads. The full impact of the third party activity, even in that context, would remain opaque to the consumer.

⁷² See Fed. Trade Comm’n, *supra* note 40; Office of the Privacy Comm’r of Canada, *supra* note 41.

⁷³ See, eg, Info Comm’r’s Office, “Guidance on the Rules on Use of Cookies and Similar Technologies” (2012) at 13, online (pdf): https://ico.org.uk/media/for-organisations/documents/1545/cookies_guidance.pdf.

⁷⁴ GDPR, *supra* note 9, arts. 13-14; “Recipients” is a broad term that includes any party who received data, whether they are a “third party” for the purposes of the GDPR or not.

⁷⁵ Office of the Privacy Comm’r of Canada, “Consent Guidelines” *supra* note 40.

These increasingly complex relationships within the data ecosystem demand regulatory responses that go beyond individual consent. We do not take up that issue here but point out that even if the law develops in this direction it will still require transparency in data flows. This transparency will be in service of meaningful accountability rather than meaningful consent. And that still leaves open the question that we are concerned with here, which is the extent to which automating privacy policies can help to bring about meaningful accountability.

Given our AppTrans study findings regarding third party code being responsible for the majority of undeclared data collection that we detected, and given these general concerns regarding the role of third parties and the requirements that mobile applications disclose the role of third parties in their privacy policies, we determined that it was important to try to focus the question of automating privacy policies on the issue of third parties. As we outline in the following section, the results were disappointing and point to some general problems in moving forward with creating usable privacy policies that can contribute to meaningful accountability.

3.3 Refining Privacy Policy Analysis Regarding Third Parties

Based on our initial experience with AppTrans for detecting non-compliance with respect to information collection, we turned to the problem of more precisely detecting non-compliance with information flows to third parties. One question we had was whether third party collection might actually be described in privacy policies through language of sharing rather than collection. In other words, we wondered whether what may have appeared in our AppTrans study to be a failure to disclose data collection might turn out to be a different kind of disclosure about third party information sharing that would not have been detected in our original study.

Rather than immediately beginning to use Amazon Mechanical Turk workers, we first conducted an experiment where we had law students with experience with privacy policies label text segments from privacy policies that described third party information flows. Each text segment was extracted out of a policy segment by searching for a sequence of sentences that use a set of keywords related to third party sharing.⁷⁶ We then divided the segment into two datasets and had each segment labeled independently by three law students. We asked the law students to label each segment by indicating (a) what type of information is shared (the students could indicate that the type of information wasn't specified); (b) whether a specific third party (i.e. Google), a general party (i.e. Advertisers) or no specific party was specified for the sharing; and (c) whether a purpose was specified for the sharing (i.e. yes or no).

The results showed that the privacy policy text segments associated with third party sharing have a high degree of ambiguity, which can result in different interpretations of the same text by different labelers. To gain an objective measure of this ambiguity, we measured how consistently segments are labeled by different labelers. For example, across over 3000 policy segments about collection labeled by MTurk workers in our AppTrans study, we found

⁷⁶ Specifically, we used share, sharing, 3rd party(s), third party(s), provider(s), another party/company, other parties/companies.

that 56% of the text segments had at least 4/5 labelers agree on the meaning of the text. It is difficult to compare this rate directly with the rate on the labels made by the law students on third party sharing since we expect law students to be better at labeling privacy policy text segments, but we also were only able to have three law students label each segment. Because of these factors, we required all three law students to agree on a label for it to be deemed an acceptable label. With this 3/3 standard, the inter-labeler agreement rates were lower across the three questions concerning third party sharing: 22% and 8% for question (a), 42% and 50% for question (b) and 46% and 58% for question (c). While empirically this suggests that there is more ambiguity in third party sharing text, we were also able to discuss the students' labeling decisions with them and obtain anecdotal reasoning for some of the discrepancies.

In general, some of the text that was problematic and led to diverging answers included vague and imprecise language that could be open to interpretation. As an example, consider:

We request information from you on our event registration or order forms. Here you must provide contact information (such as name and shipping address) and financial information (such as credit card number and expiration date). This information is used for billing purposes and to fill your orders. If we have trouble processing your order, this contact information is used to get in touch with you. We may use an outside shipping company to ship orders, and a credit card processing company to bill users for goods and services. These companies do not retain, share, store or use personally identifiable information for any secondary purposes.

There is an implication that there is information shared with a shipping company, and clearly some of it should include the types of information specified in the first sentence (name and shipping address), but it is unclear whether it will include credit card information. Similarly, it is unclear whether the shipping address will be shared with the credit card processing company. Because a disclosure involves one or more potential parties, this necessarily complicates the flow of information, leading to more opportunities for imprecise text (whether intentional or not) that can have multiple interpretations.

In other cases, exceptionally broad makes it difficult to have consistent interpretations of text. For example (within the same policy):

We may combine the information you submit under your account, with information from other [App Name Redacted] services or third parties to provide you with a better experience and to improve the quality of our services. For certain services, we may allow you to opt out of combining such information...

When you access the [App Name Redacted], we may collect your Unique Device ID ('UDID') and/or IP address and/or GPS location. We use this information to provide a tailored experience for you. The information is collected in order to determine the aggregate number of unique devices using our service or parts of our service, to track total usage, analyze data, and communicate with you more effectively. We may combine this information with information from third parties to provide you with a better experience and to

improve the quality of our service. We do not share any personally identifiable information with third parties in association with your UDID or GPS location without your explicit permission.

In both these cases, it is suggested that information is combined with third party information, but the term “combining” cannot be unequivocally distilled into whether information flows to the third party or from the third party. In addition, while no purpose is specified in the first paragraph, the second paragraph appears to specify some specific purposes. However, one could interpret those purposes to only pertain to use of personal information by the first party. In both cases, the only clearly attributable purpose of combining the information with the third party is just to “improve the quality of our service(s)”. In contrast, ambiguities that caused diverging answers in collection were less varied and the flow being described was considerably simpler as it was only between the user and the application. For example, a frequent cause of ambiguity with collection was that a single privacy policy covered both a mobile application and non-mobile application services (a web page for example), so it was unclear whether the policy was referring to the mobile app or to the other service.

Overall, we found the results discouraging. Because of the ambiguity in the policies, it would appear that there would need to be a high amount of training and coordination among labelers to ensure that privacy policy text about third party sharing will be interpreted in a consistent way (other than consistent answers that it is unclear). Moreover, it also suggests that in practice, regular mobile application users would likely have diverging interpretations of the privacy policies, thus decreasing their value as a method of obtaining meaningful consent from users. Moreover, if the end-goal is to train a machine learning model, such models need consistently labeled training sets to obtain good results, and ambiguity will naturally lead to very similar text being labeled differently because it is open to interpretation. As a result, we conclude that while such a model will predict some sort of label, if people cannot agree consistently on a correct label, then any label that is predicted will have limited value.

Our findings about the difficulties in classifying privacy policies regarding third party data sharing practices are consistent with other work in the field. Reidenberg et al’s study that analyzed how different types of users interpret policies also found higher levels of agreement in relation to data collection than in relation to data sharing.⁷⁷ The study looked at how three groups of users understood six different privacy policies: typical users - represented by crowdsourced workers on MTurk, knowledgeable users - made up of a group of graduate students in law, policy, and computer science, and expert users - made up of a group of law and policy scholars. The study subjects were each asked a series of nine questions about each policy, with four or five answer options for each. The typical and knowledgeable groups agreed with most of the experts’ mode answers for questions pertaining to data collection. There was more inter-group and intra-group disagreement when it came to sharing practices.

As we previously discussed, greater standardization in privacy policies might be one policy response that could enhance our ability to automate policy analysis given the deep problem of ambiguity and vagueness in the policies, at least in relation to some aspects such as third party sharing. But the questions surrounding third parties, meaningful accountability, and

⁷⁷ Joel R. Reidenberg et al., “Disagreeable Privacy Policies: Mismatches Between Meaning and Users’ Understanding” (2015) 30 Berkeley Tech L J 39.

automating policy analysis are much broader than the ones that appeared in our AppTrans study and its follow on work. In the following section we show how these themes have arisen in relation to the recent Facebook investigations following from the Cambridge Analytica scandal, and the question of Facebook’s responsibilities in relation to monitoring its third party app developers for policy compliance.

4 The Emerging Role of Technical Safeguards for Policy Compliance

In 2018 the Observer and the Guardian broke the story of Cambridge Analytica, revealing that the data analytics firm profiled voters with data harvested from millions of Facebook profiles.⁷⁸ This harvesting occurred through an app -- thisisyourdigitallife -- essentially a personality test created by Cambridge researcher Aleksandr Kogan for the purposes of academic research and then shared with Cambridge Analytica, in violation of Facebook’s platform policies governing app developers. In the wake of this scandal, regulators in the UK, the US, and Canada investigated Facebook’s role and they all came to two main conclusions. First, Facebook failed to adequately obtain consent from its users for its disclosures to app developers through its API. Second, Facebook failed to put in place adequate safeguards to ensure that its app developers adhered to its platform policies. In this section we look at these conclusions in light of our experience with AppTrans and outline the potential benefits as well as potential limits of utilizing technologies that automate privacy policy (or platform policy) analysis in order to fulfil these obligations involving consent and safeguards.

Facebook’s failure to obtain proper consent from its users involved a number of problematic practices. The consent failures can be broken into two main categories, the second of which is most relevant to the question of the role of policy automation: 1) Facebook’s obligation to ensure that users sharing information with “friends” on the platform understood that this would include sharing information with the apps those friends downloaded and used; 2) Facebook’s obligation to ensure that its third party apps obtained adequate consent for their access to Facebook user information. The first consent failure, regarding Facebook’s own consent practices, was also flagged in earlier investigations, such as in 2009 by the Office of the Privacy Commissioner (OPC) of Canada and again in 2011 by the FTC in the US.⁷⁹ The second consent failure, regarding the consent practices of third party apps, was more clearly highlighted in the recent investigations, although differently by the different regulators. The Canadian regulators framed the issue in terms of Facebook disclosing user information to the third party app developers, which required consent. Facebook can delegate its consent obligation to the app developers, but it is required to engage in active monitoring to ascertain whether they were in fact obtaining appropriate consent. Facebook’s practice was to require third party apps to

⁷⁸ Carole Cadwalladr & Emma Graham-Harrison, “Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach” *The Guardian* (17 March 2018), online: <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>.

⁷⁹ *In the Matter of Facebook, Inc.*, FTC File No. 092 3184, No. C-4365 (F.T.C. 29 Nov 2011) at 5, online (pdf): <https://www.ftc.gov/sites/default/files/documents/cases/2011/11/111129facebookagree.pdf>.

provide a link to a privacy policy and Facebook would monitor to see if this link was active, which was inadequate. The OPC stated:

*without ever verifying that the link actually leads to a privacy policy that explains the purposes for which the individual's personal information will be used, Facebook is not making a reasonable effort to ensure that individuals are receiving the information they need to support their meaningful consent.*⁸⁰

Going forward, therefore, Facebook needs to actively monitor the *content* of privacy policies of its third party apps and not just that there is a privacy policy. The US Justice Department had also highlighted the issue of lack of monitoring of the third-party privacy policies beyond checking for an active link to the policy.⁸¹ However, it considered this part of Facebook's general lack of monitoring for compliance with Facebook platform policies and, as we discuss below, the FTC ordered more active monitoring of this compliance.

The second main conclusion of these investigations was that Facebook failed to put in place adequate safeguards when sharing information with third party app developers through its Application Programming Interface (API) – which is the software interface that allows app developers to access Facebook user data. The language of “safeguards” is from the data protection regimes in Europe and Canada. Within the US, the FTC framed the issue in terms of the requirement to have a “privacy program”.⁸² All came to similar conclusions regarding the failure of Facebook to adequately monitor compliance with its platform policies.

Facebook relied primarily on contractual safeguards -- its set of platform policies -- to restrict how app developers could use the Facebook user information they accessed. Facebook's evidence regarding how it monitored for compliance included the use of automated tools to check whether third party apps had live links to a privacy policy, manual review of popular apps or those with high numbers of negative reviews or other indications of problems such as high numbers of deletions, and review of apps based on reports of problems from users, employees, or others.⁸³

The main problem in the Cambridge Analytica scandal was that Kogan had disclosed the data that he had collected through his app to Cambridge Analytica, in violation of Facebook's platform policies which explicitly prohibited use for commercial purposes.⁸⁴ The problem was that Facebook was unaware of the violations of its policies until an article was published in the Guardian, at which point it terminated access rights to the data and began an

⁸⁰ Office of the Privacy Comm'r of Canada, Joint Investigation of Facebook, Inc. by the Privacy Commissioner of Canada and the Information and Privacy Commissioner for British Columbia, PIPEDA Report of Findings #2019-002 (OPC, 25 April 2019), online: <https://www.priv.gc.ca/en/opc-actions-and-decisions/investigations/investigations-into-businesses/2019/pipeda-2019-002/> [hereinafter OPC Joint Investigation].

⁸¹ *United States v. Facebook, Inc.*, No. 19-cv-2184 (DDC, 24 July 2019) at paras. 118-119, online: https://www.ftc.gov/system/files/documents/cases/182_3109_facebook_complaint_filed_7-24-19.pdf [hereinafter *FTC Settlement Order 2019*].

⁸² *Ibid*

⁸³ OPC Joint Investigation, *supra* note 52, para 126.

⁸⁴ *Ibid* at para 34.

investigation.⁸⁵ What should Facebook have done? The UK ICO held that Facebook should have reviewed the terms and conditions of the app in order to determine whether these were consistent with Facebook’s policies, as well as to monitor whether the app was operating in a manner consistent with these policies.⁸⁶ What the ICO contemplates, therefore, is that contractual safeguards for sharing data with third parties are insufficient and there need to be technical safeguards. These require a means of reviewing the privacy policies or terms and conditions of the third parties and assess them against their own contractual safeguards.

The Canadian and US regulators came to similar conclusions regarding the need for Facebook to engage in proactive monitoring of compliance with its platform policies. The OPC pointed to the inherent problems in relying upon app users to raise flags about problematic data use when they had no knowledge that the app even had access to their data. It is Facebook that “knows precisely which apps get what data and when, and has the unique ability to monitor apps proactively to protect users before any unauthorized disclosure occurs”.⁸⁷ The FTC order requires, among other things, that Facebook have its third parties self-certify that they are complying with platform policies, that it deny access to third parties who fail to self-certify, that it monitor for compliance, and that it enforce violations.⁸⁸ In relation to monitoring, the FTC indicates that measures should include “ongoing manual reviews and automated scans, and regular assessments, audits, or technical and operational testing at least once every twelve (12) months”.⁸⁹

There are many important questions that we should ask about the rise of what Van Loo calls the emergence of “enforcer firms”, or the regulatory conscription of private firms to perform public regulatory duties to oversee third parties.⁹⁰ However, following from these investigations and their conclusions, we see several specific questions regarding the use of technologies to automate policy analysis. First, to what extent can we automate the analysis of privacy policies, and other platform policies, in order to facilitate the kind of proactive monitoring of compliance called for in these investigations? Second, connecting this back to our AppTrans investigation, do the conclusions of these investigations also suggest that third party advertisers who advertise through mobile apps need to actively monitor mobile apps for compliance with their platform policies, including requirements for consent?

In relation to the first question, our conclusions offer some support for the creation of technical safeguards in this context, but also some cautions. Several regulators were critical of the fact that Facebook only looked to whether an app had a link to an active privacy policy and did nothing further to determine whether the privacy policies were adequate. Our AppTrans study suggests that it is possible to create proactive systems that could determine whether an app’s privacy policy declared its data collection in relation to Facebook data. It

⁸⁵ Information Commissioner’s Office, Monetary Penalty Notice para. 43 (24 Oct 2018), online: <https://ico.org.uk/media/action-weve-taken/mpns/2260051/r-facebook-mpn-20181024.pdf> [hereinafter *ICO Penalty Notice*].

⁸⁶ *Ibid* at para 53.

⁸⁷ OPC Joint Investigation, *supra* note 52, at para. 158.

⁸⁸ FTC Settlement Order 2019, *supra* note 54, at §VII ¶E.

⁸⁹ *Ibid* at §VII ¶E.c.

⁹⁰ Rory Van Loo, “The New Gatekeepers: Private Firms as Public Enforcers” (2020) 106 *Va L Rev* 467.

would be more difficult to create a method for monitoring whether an app's privacy policy properly indicated its data sharing practices given the problems we have flagged regarding the clarity of policy language. However, a platform like Facebook is also in a position to require that its app developers have a privacy policy that is standardized in a way that facilitates automation.

The other area for policy automation arising out of the Facebook investigations is regarding Facebook's platform policies and its obligation to monitor third parties for compliance with these policies. The machine learning techniques that we used for privacy policy automation would not work in this context because one could likely not create the large database of policies required to make it work. There is still a need for some form of automation of the policies that could then facilitate the kind of technical auditing that seems called for in light of the Facebook investigations. It might be that different forms of standardization, could also apply in this context or that quite different technical solutions are required.

The Facebook investigations could also have an impact on the issue we identified through our AppTrans study, which is that a significant number of mobile apps do not disclose information collection associated with third party libraries. One way to address this problem is to focus on the app developers and their obligations and practices. However, in this context it is the third parties who are most often larger and more sophisticated than the app developers. As we outlined previously, within the data protection law framework, these third parties are often collecting personal information and must provide notice of this and obtain consent. They can delegate this to the app developers but the Facebook investigations suggest that they need to actively monitor whether the app developers comply.

Take the example of Google. An app developer who wishes to monetize her app through ads might incorporate software code from Google (the Google Mobile Ads SDK) in order to use the services of Google's AdMob. The documentation for AdMob includes the statement:

your app's privacy policy may need to be updated to reflect the use of personalized advertising (formerly known as interest-based advertising) served via the Google Mobile Ads SDK. Please take a moment to review your app's privacy policies and ensure that they are up-to-date. Because publisher pages and laws vary across countries, we're unable to suggest specific privacy policy language.⁹¹

Similarly, if an app wants to make use of Google Analytics, it must follow a set of Google policies that include the requirement to get consent in some circumstances.⁹² In addition, in order to have their app offered through Google Play, app developers have to comply with the Google Play policies, which include the requirement of having a comprehensive privacy policy.

Google therefore currently relies upon a set of policies to require that apps that make use of its services to get the required consent. This suggests two things: first, that Google could use policy-automation tools to actively monitor the adequacy of app privacy policies; and second, regulators could address the problem of third party data collection in this context by

⁹¹ See Google, "AdMob & AdSense programme policies" online: *Google Admob Help* https://support.google.com/admob/answer/2753860#Interest_based.

⁹² See Google, "Measurement Protocol, SDK, and User ID Feature Policy" online: *Google Analytics* <https://developers.google.com/analytics/devguides/collection/android/v4/policy>.

investigating third parties and not the apps. With respect to the first, this would involve the same set of limitations flagged earlier. With respect to the second, this would suggest the potential for developing a different kind of tool than AppTrans. Such a tool could look at the source code of apps in order to determine which third parties are being used, segment the apps that use the same third parties (e.g. all apps that use Google Mobile Ads SDK) and then determine whether their privacy policies are consistent with Google's obligations regarding consent. If they are not they the regulators could decide to investigate Google, rather than these apps.

The Facebook investigations are notable in the strong and consistent message they give regarding the limitations of contractual safeguards when sharing data. Sharing data between organizations and protecting that data from unauthorized uses through platform policies and data sharing agreements is a central component of our data ecosystem. It highlights how the problems of data transparency and third parties is quite complex and goes beyond the issues of third party code that we found in our AppTrans study. It also highlights how this issue will not be resolved through focusing on consumer consent alone. We need methods of ensuring meaningful accountability and these methods need to include both automating policy analysis and determining, through forms of auditing, whether data practices accord with policy authorizations and declarations. This is needed for regulators but also for organizations who want to share data and fulfil their obligations to properly safeguard it.

5 Conclusions

In recent years there has been a great deal of important work on improving the usability of privacy policies. In this paper we have argued that automating privacy policy analysis, through machine learning techniques, is a promising means of improving these policies as a mechanism of transparency. However, we have also argued that what is needed is a shift away from focusing on the role of privacy policies as enabling meaningful consent and towards focusing on their role in enabling meaningful accountability. This entails less emphasis on the role of consumer understanding and more emphasis on the role of regulators. It also entails less emphasis on creating tools like visualizations for consumer education and more emphasis on integrating policy analysis with auditing tools for regulators to determine compliance. This also entails that debates about the greater standardization of privacy policies should focus on how standardization might facilitate forms of policy automation and not on how it can improve consumer understanding in the first instance. In other words, we need to stop thinking about privacy policies in terms of whether consumers read and understand them and instead treat them as self-reporting mechanisms for data practices and then regulate them in the way we do other important disclosures such as financial disclosures or tax reporting -- by imposing standards and auditing for compliance.

Automation will enhance the usability of privacy policies as mechanisms of meaningful accountability. However, automation is no silver bullet. In our AppTrans study we found that automating policy analysis can be accurate in relation to data collection but we subsequently encountered more difficulties in relation to data sharing. It is possible to get labelers to agree that privacy policy disclosures regarding data sharing are unclear, but we are not confident that we can get labelers to agree on anything more specific regarding those disclosures. Without this, its role is much diminished. At the same time, the need for technical tools to assist in policy

analysis (whether privacy policies, platform policies, or other related policies) and in determining where data practices are consistent with policy requirements continues to grow. This was an important aspect of the requirements that came out of Facebook investigations associated with the Cambridge Analytica Scandal. Greater standardization of policy language will greatly assist in the project of automating analysis. However, as we also discussed, we need to do more to create technical environments that facilitate data audits of various sorts.