The *TL;DR Charter*: Speculatively Demystifying Privacy Policy Documents and Terms Agreements

LINDAH KOTUT, Information School, University of Washington, USA
D. SCOTT McCRICKARD, Department of Computer Science, Virginia Tech, USA

Privacy policy and term agreement documents are considered the gateway for software adoption and use. The documents provide a means for the provider to outline expectations of the software use, and also provide an often-separate document outlining how user data is collected, stored, and used–including if it is shared with other parties. A user agreeing with the terms, assumes that they have a full understanding the terms of the agreement and have provided consent. Often however, users do not read the documents because they are long and full of legalistic and inconsistent language, are regularly amended, and may not disclose all the details on what is done to the user data. Enforcing compliance and ensuring user consent have been persistent challenges to policy makers and privacy researchers. This design fiction puts forward an alternate reality and presents a policy-based approach to fording the consent gap with the *TL;DR* Charter: an agreement governing the parties involved by harnessing the power of formal governments, industry, and other stakeholders, and taking users expectation of privacy into account. The Charter allows us as researchers to examine the implications on trust, decision-making, consent, accountability and the impact of future technologies.

CCS Concepts: • **Security and privacy** → **Human and societal aspects of security and privacy**; • **Human-centered computing** → *Human computer interaction (HCI).*

Additional Key Words and Phrases: speculative design; consent; privacy; privacy policy documents; policy simplifications; text summarization; GDPR; design fiction

ACM Reference Format:
Lindah Kotut and D. Scott McCrickard. 2022. The *TL;DR Charter*: Speculatively Demystifying Privacy Policy Documents and Terms Agreements. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP, Article 23 (January 2022), 14 pages. https://doi.org/10.1145/3492842

1 PREAMBLE

Documents on expectation of use have been a staple since software could be directly downloaded and used by consumers/users (we use both terms interchangeably in this document). These agreement documents outline the terms of use, set expectations on software performance, and protect the software developer(s) against indemnity for unexpected harm caused by software use. The user agreeing to terms, follows an all-or-nothing principle: usage of the software implies consent to all the terms of use and privacy policy outlined by the developer. The terms of use documents have become standard beyond desktop software and into applications downloaded for handheld devices and then to other devices classed under Internet of Things (IoT)–including stock applications bundled with those devices.

Authors' addresses: Lindah Kotut, kotut@uw.edu, Information School, University of Washington, USA, Seattle, Washington, 43017-6221; D. Scott McCrickard, mccricks@cs.vt.edu, Department of Computer Science, Virginia Tech, USA, Blacksburg, Virginia, 24060.

2573-0142/2022/1-ART23
https://doi.org/10.1145/3492842

In the age of targeted advertising, privacy policy documents are often included as a separate document that outlines how a user's information will be collected, stored, used, and shared with third parties. Given the number of applications an average smartphone user has on their phones, the number of websites they visit, and the increased ubiquity of IoT devices they have in their homes, these contracts are useful in providing a means for personal curation of what has/will be collected and stored by first-parties, and what will be shared to third-parties–together with the implications that this has on the data. However, the terms of use and privacy policy documents (we will refer to both as contracts) tend to be long, requiring advanced reading skills and/or an understanding of legalese [35]. Crucially, not all companies abide by the contracts that they share with the user, and some do not disclose all of the data that will be collected and what is done with the collected data [27]. This combination of factors make it difficult for a user to be able to give informed consent.

Myriad stakeholder recommendations exist on what should be done to ease user decision making: automating the reading of contracts and providing recommendations by using privacy assistants [4], leveraging visual representation using dashboards [28], and using crowd-sourced approach to summarize the contracts [25]. These strategies while useful, are subject to lost interest and limited participation [33], do not compel the developers to be transparent, and either hinder, or do not ease the process of consent. Laws have the power to compel–and have been used to enforce the preservation of personal privacy regarding health data [23], or personal data writ large [30]. However, these laws tend to be fragmented, and are often not interoperable–as they depend on country or regional context. This renders making use of the strengths of proposed approaches difficult.

To address these limitations, we leverage design fiction to introduce the *TL;DR* Charter: a document that specifies the rights of world-wide signatories by providing guidance on data handling, enforcing accountability, and notably, giving the developers and other producers the flexibility of requesting or requiring access to data they need, and empowering the users to make informed decisions about whether to accept the terms. Importantly, the Charter provides a unified system that is universal in its very nature–enfolding multiple, often conflicting laws from different countries and regions. This affords the user the opportunity to exercise informed consent using a combination of at-a-glance features, summarized features, with an additional option to read the entire contract spelling out the terms of use and privacy expectations.

2 BACKGROUND

Before presenting the Charter, we first outline challenges to user informed consent. We include the recommendations and limitations as outlined by various stakeholders who have considered different avenues to aid consumer understanding and producer compliance.

2.1 The Balkanization of Policy Agreements

Terms of use and privacy policy documents (contracts) are intended to outline expected use and privacy expectations. As the amount of data that could be collected have increased, and companies have grown to include subsidiaries, the contractual documents have grown in tandem–in an attempt to cover as many eventualities as possible. Often however, the contract language tend to hedge, obfuscate, and/or downplay risks [35]. The contract length and the sheer number that are produced make it impractical for the consumer to monitor, causing cognitive fatigue [26]. The frequency of contract updates and the obfuscated language place further burdens on the consumer to determine recourse and protective measures to undertake in case of a breach [35], with no confidence that their choices will be honored [12]. Some providers also leverage dark patterns that seek to nudge

users towards clicking "Accept" without reading the agreements [31], often presenting asymmetric and incomplete information that erode the process of decision making. [1][1]

To ford this trust gap, users tend to resort to other measures provided by platforms (for example the number of permissions requested by a smartphone application) and browsers [19], as interim (and incomplete) replacement signals on what the contracts stipulate–to claw back some agency over their personal data. Other visual approaches have included browser add-ons that provide a graphical representation of which parties have access to the user data [19]. These approaches, while useful, are neither sustainable, nor scalable–tending to be casualties of lost interest and participation [33]. In addition, while summarization are useful at-a-glance services, they are reactive [31], tend to interfere with user consent [10], and do not even enforce compliance.

Human-in-the-loop approaches have been used to provide a measure of scalability both in effecting large-scale labelling: using crowd-sourcing to provide at-a-glance determination of expectations [25], and in leveraging expert assessment to present layered feedback to the user [13]. These attempts to counter notice-and-choice approaches are also meant to pressure providers towards transparency and compliance with laws [6]. Yet they still suffer from the scalability problems, and become stale once a contract is updated. Other approaches geared towards older consumers who use IoT devices propose the use of personalized privacy assistants [4] and privacy profiles [18]. This is done to simplify the contracts by only providing relevant/customized recommendations, including the use of dashboards [28]. These approaches do not effectively inspire trust. For example, older users still tend to rely on friends and family members over automated means to elicit trust measures [20]. Expert advice also tend to be a legion: making it difficult for the consumer to prioritize which one to act upon [24]: highlighting the need for a holistic solution.

The disclosure deficiencies tend to paint providers with a broad villainous brush. However, there is difficulty in ensuring compliance due to language ambiguity. Developers need help with disclosing the third-party libraries they use, and also streamline enforcement [34]–which require a lot of resources [26]. Proposed assistive approaches have included compliance management platforms as a way to effectively increase compliance [22]. Knowing the limitations (scalability, trust signals, variability of documents, trust enforcement), provide further opportunities to consider higher-level approaches to address these limitations instead of considering them at the application level. This need provide a space to leverage a speculative design approach that would consider unification of the varied approaches and solutions.

2.2 The Letter and the Spirit of the Law(s)

While policy and user agreements provide an opportunity for informed consent, the burden falls on the consumer in case of harms, often with little to no recourse. Laws are the avenues to enforce compliance, having the power to compel actions and to penalize infractions. Existing laws include those that govern how practitioners in the United States handle health data through the Health Insurance Portability and Accountability Act (HIPAA) [23]. Different states within the United States, and various countries also have laws governing different aspects of user data [11]. These laws are sporadic, and are rarely updated: HIPAA for example had its last major update in 2013[2]. The implication of the lack of maneuverability with changing times and changing technology results in ineffective enforcement. An example regards the number of days mandated for a user to be notified in case of a breach: the law requires that the user be informed through physical mail via the post

[1]Researchers make use of the "market for lemons" metaphor to typify how these dark patterns mimic what is done in the car sale industry to sell vehicles with manufacturing defects.
[2]https://www.hhs.gov/hipaa/for-professionals/index.html

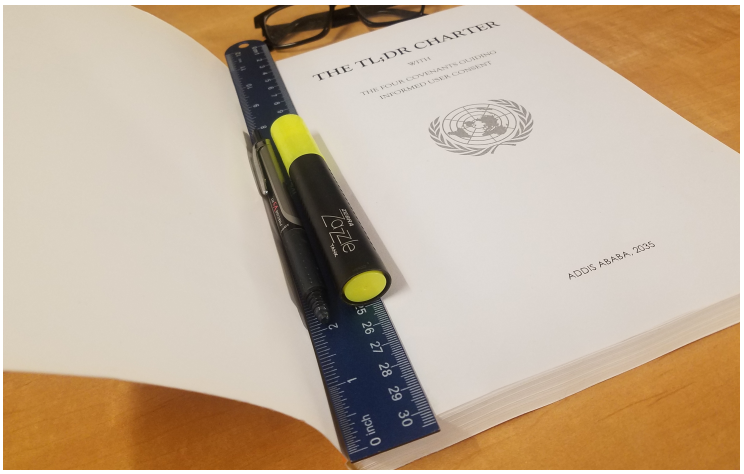

Fig. 1. The *TL;DR Charter* is aimed at streamlining the process of user consent. It was ratified in 2035.

office [35], depriving them of the opportunity to promptly address possible harms caused by the breach.

The European Union's General Data Protection Regulation (GDPR) [30] offer an example of an approach to writing laws that are comprehensive, and involve multiple countries as signatories. GDPR addresses both the privacy expectations and proper enforcement. When compared to existing laws, the timeliness of updates–in keeping with changing technology, and its core foundation/tenet in consent, it is considered the current gold standard. Researchers using the GDPR as a standard to assess compliance to the law have found that while there was some observed reduction on the amount of tracking since the law went into effect, this did not translate generally [27]. In fact, while websites gave users a glanceable means of determining compliance to GDPR via "levelled choices" on their tracked activities [14], the websites tracked the users anyways–before, and even after the user has stated their preferences against being tracked. This direct contravention of both the letter and the spirit of the GDPR laws was observed to be especially prevalent in news sites whose revenues are primarily based on advertising [8]. The number and type of third-party sites also tended to change, making it difficult to keep track either manually or automatically, with no history of the providers and changes available to the users. All these actions reveal the challenges that still beleaguer the decision making process on the part of the user, and the contravention to the spirit of the GDPR agreement (which was expected to have resulted in less incidences of third-party tracking) [27].

These abiding challenges across user agreements and laws highlight the need for a different approach. We outline the *TL;DR* Charter in the following section. The Charter is geared towards presenting a unified enforcement approach that is easy to follow on both the part of the provider and consumer, easy to notify in case of a breach, and easy to indemnify in case of harms.

3 THE TL;DR CHARTER

The *TL;DR*[3] *Charter*[4] was ratified in 2035 after signatories–made up of non-governmental organizations (researchers included), companies, and governments, agreed on a framework to ease

[3]Too long; didn't read

[4]We define a charter as a "A grant or guarantee of rights, franchises, or privileges from the sovereign power of a state or country": https://www.merriam-webster.com/dictionary/charter. Inspired by 'The UN Charter' (1945).

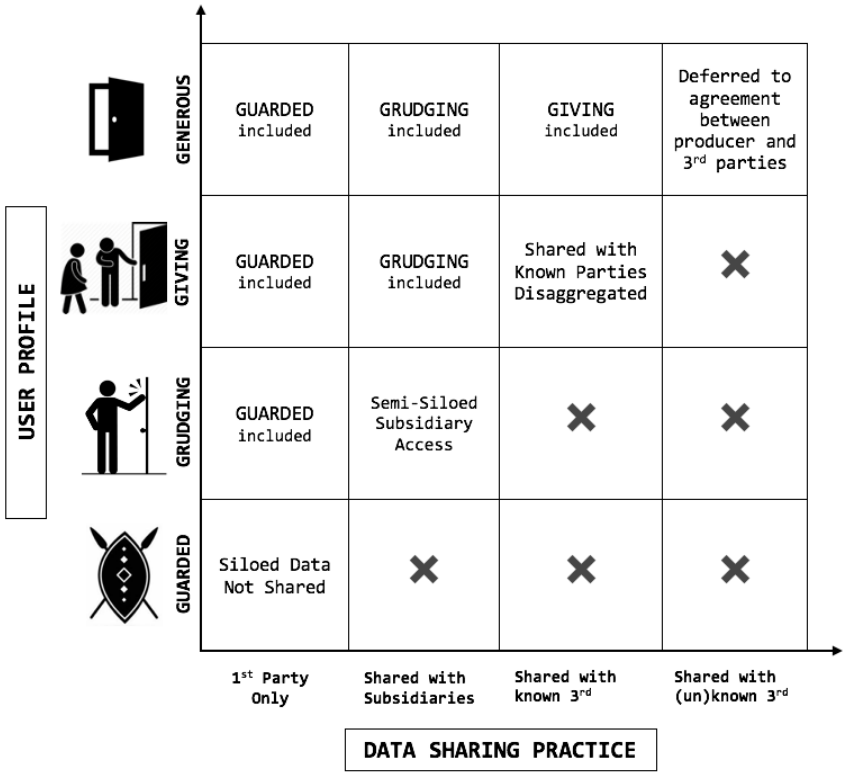

Fig. 2. User profiles based on data sharing preferences matches the covenants described in the *TL;DR Charter*. They provide enforceable guidelines on data collection, storage and sharing. There is space for flexibility and negotiation within the parameters of each user profile that preserve trust between and among parties.

the consent process (see Figure 1). The process of collecting the signatures began as a grassroots mobilisation in 2026. It took five years for the original draft to be completed, and a further four years for the signatories to agree on the final draft. Therefore while the *TL;DR Charter* is currently at it's infancy, it is important to describe the benefits it provides to both the users and the providers. The Charter further eases the process of enforcement–the prerogative of the government signatories, yet allowing the space for additional country-specific laws and interpretation of recourse actions.

In its essence, the Charter provides a way for the signatories to have one place to combine their strengths. This is leveraging lessons from GDPR adoption and enforcement, in addition to other community-specified standards[5]. The Charter also considered and adopted some user input through a comment period that preceded the Charter's ratification. Additional researchers feedback provided an understanding of best practices to adopt.

3.1 Charter Purpose Specifications

The Charter is first and foremost, community focused. While current enforcement strategies are specific to certain aspects of privacy (e.g. governing how personal data is collected and used, and addressing some power differentials for example giving control of personal information back to the user when considering the "right to be forgotten"), the *TL;DR Charter* allows for more flexibility.

[5]see: https://www.w3.org/Consortium/

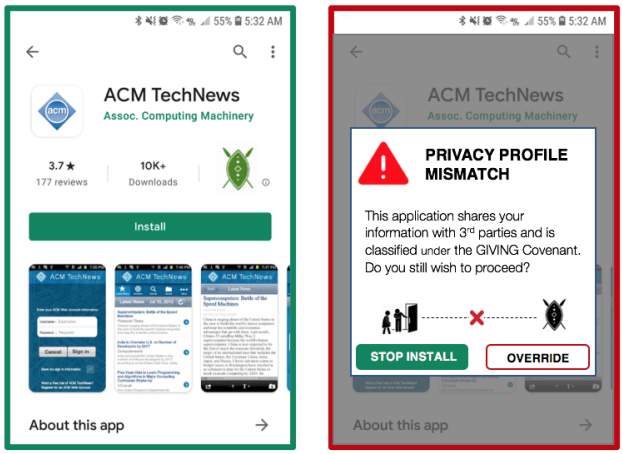

Fig. 3. Charter enforcement would result in at-a-glance application using both border color and covenant icons as guides. The left image is an example of a PlayStore application that match a user's preferred covenant. Right screen shows triggered prompt when a user begins a download of an application that is more permissive than their profile allows.

The use of covenants (defined as formal agreements) describes frameworks that allow for better understanding of the trade-offs between what is gained and what is given in return.

The covenants are divided into four broad groups geared towards consent categories. Figure 2 provides a visual representation of the user profiles modelled after their expectation of privacy. In the following subsections, we will introduce each of the four covenants: describing the framework, the interaction between signatories, and expected enforcement principles.

3.2 Guarded Covenant

The most restrictive of the four covenants is the *Guarded Covenant*. Users choosing this profile will be guaranteed that their information if collected, will only be used by the first party in the agreement within the bounds of the service provided. Data access to subsidiaries are excluded in this covenant, with the choice of more permissive privacy profile being at the user's discretion. The provider (of technology, tool and/or service) may not impose that choice on the consumer.

The choice of the *Guarded Covenant* has implications on what technology is available to the user. For providers of who rely upon collecting sensitive information from users in order to provide services such as health tracking, this covenant in advantageous to both parties. The boundaries surrounding how data is handled in this covenant engenders and enforces trust. From the user's perspective, there are providers who can/should be forced into providing service in this tier in order to enforce trust and protect user data. Examples of these are providers of institutional Support who are impelled by this covenant to ensure that there is equitable access to fundamental resources. Students using a learning platform for example, and preferring this covenant, should not be shut out of the learning materials required for their edification. In other aspects, users will likely be expected to pay for premium access as a compensation to the provider on potential revenue lost on target advertising – following the "willing buyer; willing seller" principle [16].

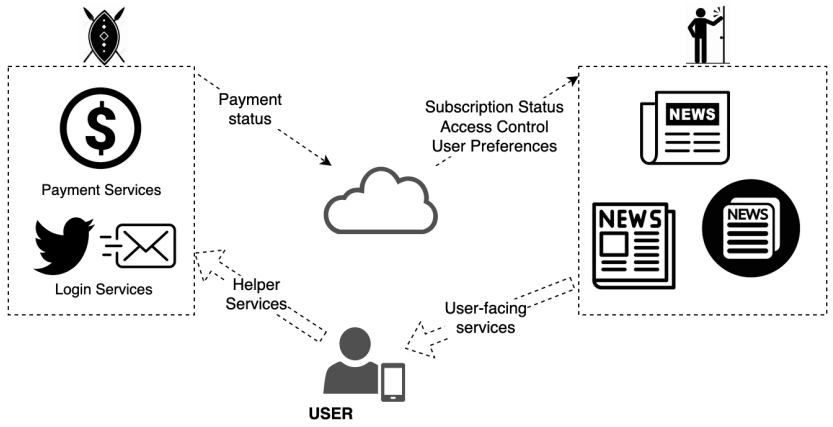

Fig. 4. A paid news subscription service can share users information amongst the newspaper subsidiaries, while also leveraging third-party services classified under the Guarded covenant to provide any additional helper services.

3.3 Grudging Covenant

The *Grudging Covenant* provide for a means of personal information to be shared laterally with the provider's subsidiaries. Say a user signs up to use an email platform; the platform provider may use that user information in the music platform offered by their umbrella. The user's information is expected to be siloed under the same umbrella agreement, and should not be shared by third parties, with acceptable caveats.

This covenant can be compared to the principles of HIPAA, that guide the handling of personal information by "business associates" that are required to facilitate service provision to the user [23]. The "associates" are also signatories of the same covenant[6].

Most subscription services would be classified under this covenant. The transaction: money for services, enforcing this covenant. Signatories in this covenant can leverage other services prescribing the Guarded covenant to provide additional services to the user with the clear understanding about how those services are rendered and reused, as shown in Figure 4.

3.4 Giving Covenant

The *Giving Covenant* is the first covenant in the permissive spectrum in that it grants third-party data access. Data shared with such parties will expected to be anonymized and aggregated, with strict provisions requiring that the third parties be known and be signatories of either *Giving*, *Grudging* or *Guarded* covenants. In case of harms caused to the user by third parties, the indemnity burden would be carried by the responsible known party.

Users in this profile will be expected to have access to a wide varieties of services and be exposed to targeted advertising. The providers will be expected to provide explicit accounting of all third-party signatories and to ensure that each follow the spirit and the letter of the *Giving Covenant*. User information and autonomy is most at risk in this covenant.

[6]For example, Facebook owns Instagram: login credentials for one, can be used in the other. For this connection to go into effect, Facebook (and Instagram) have to be signatories of the Grudging covenant. The Guarded covenant–the most restrictive covenant, does not allow any such sharing.

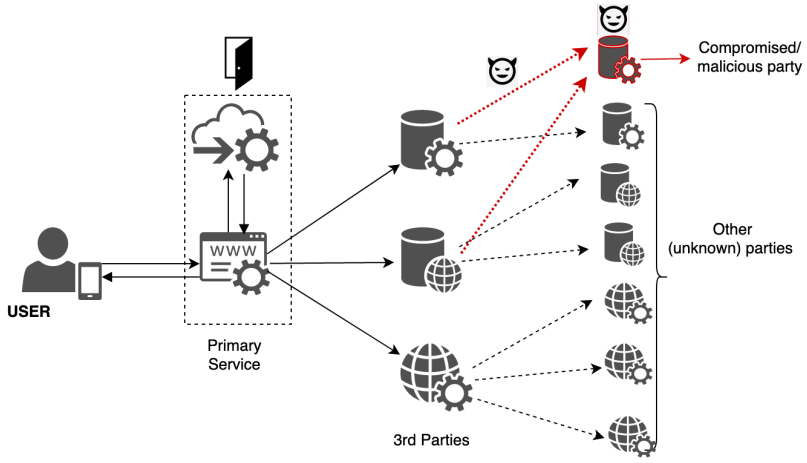

Fig. 5. *Generous* covenant reflect the risks/rewards that that a user is exposed to: Expanded service in return for access to personal/device information, but with the risk that the user will not be aware of the breach of nth parties, or data shared with deliberately malicious services. This risk abides even with the *TL;DR* Charter presence and enforcement.

3.5 Generous Covenant

Should the provider opt not to account for all the third-party who would have access to (aggregated) data, they can be signatories to the *Generous Covenant*. While the covenant requires the providers furnish reasonable protection to user data, and to account for the third parties who will have access to the said data, they are not required to ensure that these parties are signatories to peer, or more restrictive covenants. This makes it a high risk, high reward category for the user (Fig. 5).

Therefore, a provider may be implementing the *Guarded Covenant*, but wish to provide free access in return for advertising, then they would sign onto the *Generous* platform to place the onus on vetting the access to third-parties onto the user. Users who sign up to this covenant would expect to have the widest variety of technology services to choose from, given the accessibility they offer to the providers. They would as well have access to any technology product or service that are in the more restrictive continuum.

3.6 Further Guiding Principles

The use of covenants allow for a graceful transition between services should the producer wish to no longer provide their services in the given covenant tier. Producers will provide a sun-setting period allowing the users to make informed decisions on whether to sign up to a more permissive profile if that was offered, download/destroy their data, or find a replacement service. The covenants also protect the user data should the producer signatory choose to merge with other providers, or sell their service to another company outright. The users would otherwise be grandfathered, with the new purchaser mandated to preserve the same agreement provided in the user tier. If neither is possible, then the sun-setting procedure will be observed.

While the Charter outline the responsibility of the provider to alert the users on data breach in stark terms, it is up-to the laws governing the provider/user agreement to guide the recourse and compensation. However, the Charter limits the possibility of circumvention of the spirit and the letter of the agreement, and imposes trust penalties in such events, together with providing regular

amendments to account for any overseen loopholes, while also providing a means of indemnifying parties that leverage the user data in a manner not prescribed in the covenant that both parties signed on to.

3.7 Enforcing Trust

The primary objective of the *TL;DR Charter* beyond simplifying the process of compliance, is to also engender user trust. The Charter's strength is in streamlining the agreement process, and then allowing the stakeholders in these process to approach consent in a straightforward manner (see figure 3 as an example of how application download would be impacted), and more easily identify signatories that contravene the agreements.

Beyond the penalties enforced by the resident laws and allowable under the Charter, a breach of covenant will incur different trust score penalties depending upon the severity and the culpability of producer in the breach as per agreed-upon measures. There is even a possibility for other score measures to be included, for example an obfuscation score that penalizes the provider who obfuscate terms of agreement, make it difficult to summarize and understand these documents. As with the current laws, providers will be compelled to be signatories in the Charter to provide these services for their citizens. However, the ease of monitoring and verification will provide a means of consensus on the agreement regarding compliance, trust scores, and other envisioned labels depending on future needs.

4 IMPLICATIONS

After providing details regarding the provisions described in the Charter covenants, we explore the *TL;DR Charter* implications in two broad applications: (1) its capability to unify disparate technology providers to aid user consent, and (2) future proofing agreements. We elaborate on both below.

4.1 Unifying Disparate Data

The most visible implication of the *TL;DR Charter* is in unifying discordant agreements. The increased ubiquity of Internet of Things (IoT) for example highlight the disparate way that data is collected, aggregated, and shared–with the attendant difference in actions that should be taken in case of a data breach. Implications of data collected by fridges and house lights if breached, can be used to chart activities in a household, whereas a breach of data collected by security cameras can be used to supply facial recognition input. The unified framework provided by the Charter offers the answer to the different collection points, addresses different concerns raised with different associated risks at present, and unifies the different laws and contracts (if provided) that currently govern these disparate systems.

Beyond considering jurisdictions in terms of laws governing them, the Charter provides a means of simplifying access that has heretofore been limited by language fluency (of both providers and consumers), and knowledge of laws governing cross-border storage and use of collected data[7].

The Charter's provision of user-preferred profile is a first step in facilitating informed consent, while also enforcing the transaction of what is considered appropriate. For example, users in the most restrictive covenant may be required to pay for access to a technology or tool, whereas others in permissive categories would be allowed to get free access in exchange of granting access to their data and/or allowing targeted advertisements to be served to them. While this model is present in some applications in devices–the Charter allows it to permeate to other spaces.

[7]For example, US newspapers opting to deny access to users in the GDPR-governed states instead of providing GDPR-compliant versions https://www.theguardian.com/technology/2018/may/25/gdpr-us-based-news-websites-eu-internet-users-la-times

4.2 Future Proofing

The signatories to the Charter understood the far reaching consequences of having a unified framework. Beyond enforcement, it would provide new data points allowing users to 'search' for tools by their privacy profile: restricting the results to applications/websites that enforce the covenant that matches their preferences, while also including those that enforce more restrictive covenants. The Charter further provide a means for producers with no in-depth understanding of the law to have a means of creating agreement documents that preserve what they need while complying with the Charter requirements.

The Charter also defines a space where a trust rating can be envisioned, and uniformity can be cultivated–where machine learning approaches can then be used to ease the process of summarizing these documents. As it currently stands, there are no provisions guiding how user agreements and privacy policy documents should be written, and how best to advise the user as the language changes (where they have changed and what that implies regarding original understanding of the document). Industry courtesy standards exist for example in (large) providers sending emails alerting the user to the change, providing a summary of the changes, and/or displaying an additional agreement label once the agreement comes to effect. However, the user is neither allowed a choice, nor can they request their data to be forgotten in case they would no longer prefer to agree with the updated terms. Further, the length and the language of the terms favor users who have attained college-level education in western countries, and are fluent mostly in English and/or French (or sometimes German). Majority of the world languages are typically not used in these documents, impacting the large non-English/non-French speaking population

Taking into account future personal technology e.g. skin technology, DNA profiles and future assistive technology e.g. robots for hire, we anticipate that they will involve terms of use with varied granularity and governed by time and requirements that were vague at the Charter's ratification phase. The known importance of future proofing informed both the description of the four covenants to be flexible enough to allow for upcoming and even unanticipated technology needs.

Considering permitted data, while a user may have granted access to personal data that was then packaged and reused for targeted advertising the needs might change. In the next technology wave, there may be a different need for data beyond advertising. The *TL;DR Charter* provides a means for the user to have autonomy over what is used to provide personalization (for example what is offered in social media timeline) over what is used for targeted advertising–in this case, enforcing the *spirit* of the Charter. The *TL;DR Charter* can further provide an opportunity for designers to either envision new tools, or reuse old designs to aid consent: for example in graphing connection of all third parties [19].

4.3 Blind Spots

While the Charter is restrictive in the defined boundaries between the four covenants, we note that there is flexibility inherent within a specific covenant. This is important for user protection. For example, should the data provider follow the "move fast and break things" [8] approach in the race to release [a] minimum viable product(s), the covenant agreement ensures that the user is protected within the bounds of the covenant–and that they trust the enforcement of covenant breach, and indemnity in case of harms caused.

We note here that there are aspects that are not well accounted by the Charter. For example: say a homeowner advertises their house for rent, with utilities included in the rent amount. Any IoT profiles of devices in the house belong to the landlord. However, the user of these devices is the tenant–who may not be aware of the profiles governing the user-producer relationship, or

[8] A software engineering ethos attributed to Facebook: https://www.businessinsider.com/mark-zuckerberg-2010-10

the implications of the agreement. Further, data collected by these devices do not belong to the homeowner who is the signatory, but by the tenant [29]. In case of a data breach (or data used to target the owner), the injured party is the tenant, yet the indemnity would most likely be granted the landlord.

On inequalities, there are populations that cannot afford privacy. For example, users who are homeless in the United States are required to provide their personal information in order to receive social services [21]. But there is a lack of informed consent regarding the use and storage of their data, including how this data is used to make judgements about their needs and resources availed to them based on the data. This, and other similar questions of power differential and access to technology, provide a consensus point for collaboration across disciplines and jurisdictions, with the aim of ensuring that all users regardless of access, are granted the same opportunity to choose how they manage their data.

These examples bring to bear aspects of relationships that are not in the jurisdiction of the Charter, but instead are governed by local laws and/or reveal abiding societal concerns. While there are provisions guiding public institutions on how to provide a means for users' preferences to be respected, the homeowner-tenant relationship example above showcases the limitation of the Charter covenants.

As provisioned in the Charter's Guiding Principles, the *TL;DR Charter* attempts as much as possible, to mitigate power asymmetries between providers and users, and between, and among countries. This does not stop possible cases of abuse: conflicts between countries, lobbyists influencing how laws are set up, signatory countries choosing to leave and not offer options to their citizens, and third-party crowd-sourcing used for nefarious purposes. However, due to the Charter's core in providing a means of easing the consent process, the user will be poised to make the most educated decisions regarding actions to be taken in that regard.

4.4 Moving the Discussion Along

The Charter in sweeping most present concerns regarding providing tools and resources to effect user informed consent, clears the path for us to clearly perceive the delineation between the good-faith attempts at obtaining the user consent and those that misrepresent intent. The *TL;DR* Charter for its thoroughness, relies upon three prongs: (1) The service providers correctly classifying their services to the right covenant (2) the good faith summarization of privacy policy documents (3) not governing the services that provide these summarizations, the user–as a consumer and not an active partner in consenting to the categorizations as shown to them.

The three concerns allow the space to extend the discussion by merging aspects of user privacy to the ethics of the tools designed to ease the process of consent, and the assumption of good faith application. We are then able to join forces with researchers in the named disparate fields to also enfold and consider user autonomy across contexts and cultures to give granularity and depth to how we design tools and technology intended to support the user, and not undermine their autonomy.

In essence, the Charter provides a means to keep the parties honest in their intentions–in keeping with the letter/spirit of the agreement. This encompasses most small-scale providers who may not be conversant with applicable laws, especially if their user base straddle different jurisdictions. The Charter however, does not solve for parties whose intentions are to collect user data for non-documented/black-boxed, or unconsidered uses, but it serves to minimize the noise in considering (emergent) edge-cases and provide a space for marshalling of forces in considering these cases and emerging problems, and perhaps move the discussion, research, development beyond the problem of supporting user consent.

5 AUTHOR(S) STATEMENTS

Design fiction is an approach to explore possible futures [7], and has been used as design tools in various forms within HCI research depending on design need [2, 3, 15, 17, 32]. In this work, we define the *TL;DR Charter* and place it within an alternative world, following a world-building design fiction approach [5]. This approach allows us to examine the issue of privacy and consent at the macro level, while also building upon previous research leveraging design fictions to consider the ethics of data collection [9] and user agency in the collection [16].

Considering speculative futures grants us freedom to chart future research directions instead of focusing on papering over the cracks highlighted by existing loopholes. The *TL;DR Charter* is woven using three strands: (1) the use of design fiction as a tool, (2) the focus on informed consent, and (3) the consideration of laws governing privacy. As we've outlined in the Background Section, while there are examples of industry-defined and crowd-enforced standards, they tend to suffer from either lack of adoption and/or lax enforcement. GDPR laws have given a glimpse of possible successful future that involve multiple stakeholder agreements that span countries, with spelled-out penalties for breaches. The *TL;DR Charter* considers a universal agreement on privacy enforcement uniting varied jurisdictions and streamlining understanding, adoption, and enforcement, while at the same time addressing the complexity inherent in GDPR that make it difficult to understand how best to comply/how to check for compliance. Importantly, the Charter adopts an unapologetic user-centric approach in serving as an advocate and nurturing trust.

Given the speculative future, we are also able to consider the possible impact of the current research trajectories. While research designing tools to aid the users in demystifying different platforms persist, we find that this is not sustainable as more IoT applications come online. While we can envision researchers finding opportunities to provide tools to aid the understanding of these platforms, it is more difficult to envision an average user being in the know about the availability of these tools and to keep up with the changes/updates. The *TL;DR Charter* simplifies this process and frees researchers to consider how best to approach such futures: anticipating lawmakers needs and having a stake in informing future laws for example.

Drawing from research in privacy and policies, we elicited various implications of the *TL;DR Charter* in the enforcement. We hope that the macro-level approach provide a space to consider the impact at the micro-level, and the sufficiency of adopting a user-first advocacy approach. The *TL;DR* Charter is meant to spark such discussions: should we find a once-for-all solution to the process of obtaining informed consent, what is revealed, and how does this knowledge then impact the current research directions?

6 CONCLUSION

Data ownership and consent to share this information should be sacred. Privacy policy documents and terms documents are a space for producers to articulate expectations and how user data is collected, stored, used and shared. However, these documents tend to be long and convoluted: not easy to understand and tend to cause cognitive fatigue. Users thus often tend to accept the terms without reading them. Current approaches attempting to address these limitations suffer from being varied and scattered. They do not scale well, and there is a lack of unifying frameworks to guide tools and automation to help the users make decisions regarding their own privacy.

Leveraging design fiction as a tool, and inspired by successful approaches including overarching laws such as GDPR, we propose the *TL;DR Charter* as means to unify the scattered recommendations in a streamlined process. It adopts a user-advocate approach in setting expectations regarding privacy and some agency over data collected. The Charter offers this flexibility by describing four covenants that respectively describe four different privacy profiles with attendant enforcement

in case of breaches. We examine the Charter's implications including how it guides the future of research, in providing a means of categorizing the longevity of existing problems, and further, on the limitations of the Charter especially as it regards vulnerable populations.

REFERENCES

[1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM Comput. Surv.* 50, 3, Article 44 (Aug. 2017), 41 pages. https://doi.org/10.1145/3054926

[2] Eric P. S. Baumer, Mark Blythe, and Theresa Jean Tanenbaum. 2020. Evaluating Design Fiction: The Right Tool for the Job. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) *(DIS '20)*. Association for Computing Machinery, New York, NY, USA, 1901–1913. https://doi.org/10.1145/3357236.3395464

[3] Mark Blythe. 2014. Research through Design Fiction: Narrative in Real and Imaginary Abstracts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 703–712. https://doi.org/10.1145/2556288.2557098

[4] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the Design of a Personalized Privacy Assistant for the Internet of Things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376389

[5] RTD Conference, Paul Coulton, Joseph Lindley, Miriam Sturdee, and Mike Stead. 2019. Design Fiction as World Building. https://doi.org/10.6084/m9.figshare.4746964.v1

[6] Lorrie Faith Cranor. 2012. Necessary but not Sufficient: Standardized Mechanisms for Privacy Notice and Choice. *J. on Telecomm. & High Tech. L.* 10 (2012), 273.

[7] Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming.* MIT press.

[8] Steven Englehardt and Arvind Narayanan. 2016. Online Tracking: A 1-Million-Site Measurement and Analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (Vienna, Austria) *(CCS '16)*. Association for Computing Machinery, New York, NY, USA, 1388–1401. https://doi.org/10.1145/2976749.2978313

[9] Casey Fiesler. 2019. Ethical Considerations for Research Involving (Speculative) Public Data. *Proc. ACM Hum.-Comput. Interact.* 3, GROUP, Article 249 (Dec. 2019), 13 pages. https://doi.org/10.1145/3370271

[10] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. 2016. How Short Is Too Short? Implications of Length and Framing on the Effectiveness of Privacy Notices. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 321–340. https://www.usenix.org/conference/soups2016/technical-sessions/presentation/gluck

[11] Pam Greenberg. 2020. Security Breach Notification Laws. https://www.ncsl.org/research/telecommunications-and-information-technology/security-breach-notification-laws.aspx.

[12] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2020. "It's a Scavenger Hunt": Usability of Websites' Opt-Out and Data Deletion Choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376511

[13] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2019. An Empirical Analysis of Data Deletion and Opt-Out Choices on 150 Websites. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. USENIX Association, Santa Clara, CA. https://www.usenix.org/conference/soups2019/presentation/habib

[14] Xuehui Hu and Nishanth Sastry. 2019. Characterising Third Party Cookie Usage in the EU after GDPR. In *Proceedings of the 10th ACM Conference on Web Science* (Boston, Massachusetts, USA) *(WebSci '19)*. Association for Computing Machinery, New York, NY, USA, 137–141. https://doi.org/10.1145/3292522.3326039

[15] Lindah Kotut and D Scott McCrickard. 2020. Amplifying the griot: Design fiction for development as an inclusivity lens.

[16] Lindah Kotut, Timothy L. Stelter, Michael Horning, and D. Scott McCrickard. 2020. Willing Buyer, Willing Seller: Personal Data Trade as a Service. In *Companion of the 2020 ACM International Conference on Supporting Group Work* (Sanibel Island, Florida, USA) *(GROUP '20)*. Association for Computing Machinery, New York, NY, USA, 59–68. https://doi.org/10.1145/3323994.3369899

[17] Conor Linehan, Ben J. Kirman, Stuart Reeves, Mark A. Blythe, Joshua G. Tanenbaum, Audrey Desjardins, and Ron Wakkary. 2014. Alternate Endings: Using Fiction to Explore Design Futures. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI EA '14)*. Association for Computing Machinery, New York, NY, USA, 45–48. https://doi.org/10.1145/2559206.2560472

[18] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. 2016. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Proceedings of the Twelfth USENIX Conference on Usable Privacy and Security* (Denver, CO, USA) *(SOUPS '16)*. USENIX Association, USA, 27–41.

[19] Mozilla. 2019. Firefox Lightbeam Extension. https://support.mozilla.org/en-US/kb/lightbeam-extension-firefox-no-longer-supported.

[20] James Nicholson, Lynne Coventry, and Pamela Briggs. 2019. "If It's Important It Will Be A Headline": Cybersecurity Information Seeking in Older Adults. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3290605.3300579

[21] Safiya Umoja Noble. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press.

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

[23] U.S. Department of Health and Human Services: Office for Civil Rights (OCR). 2016. https://www.hhs.gov/hipaa/for-professionals/faq/business-associates/index.html.

[24] Elissa M. Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L. Mazurek. 2020. A Comprehensive Quality Evaluation of Security and Privacy Advice on the Web. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 89–108. https://www.usenix.org/conference/usenixsecurity20/presentation/redmiles

[25] Hugo Roy. 2020. Terms of Service; Didn't Read. https://tosdr.org/index.html.

[26] Sean Sirur, Jason R.C. Nurse, and Helena Webb. 2018. Are We There Yet? Understanding the Challenges Faced in Complying with the General Data Protection Regulation (GDPR). In *Proceedings of the 2nd International Workshop on Multimedia Privacy and Security* (Toronto, Canada) *(MPS '18)*. Association for Computing Machinery, New York, NY, USA, 88–95. https://doi.org/10.1145/3267357.3267368

[27] Jannick Sørensen and Sokol Kosta. 2019. Before and After GDPR: The Changes in Third Party Presence at Public and Private European Websites. In *The World Wide Web Conference* (San Francisco, CA, USA) *(WWW '19)*. Association for Computing Machinery, New York, NY, USA, 1590–1600. https://doi.org/10.1145/3308558.3313524

[28] Welderufael B. Tesfay, Peter Hofmann, Toru Nakamura, Shinsaku Kiyomoto, and Jetzabel Serna. 2018. I Read but Don't Agree: Privacy Policy Benchmarking Using Machine Learning and the EU GDPR. In *Companion Proceedings of the The Web Conference 2018* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 163–166. https://doi.org/10.1145/3184558.3186969

[29] Swapna Thorve, Lindah Kotut, and Mary Semaan. 2018. Privacy Preserving Smart Meter Data. (2018).

[30] The European Union. 2016. General Data Protection Regulation. https://op.europa.eu/s/ocJq.

[31] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)Informed Consent: Studying GDPR Consent Notices in the Field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 973–990. https://doi.org/10.1145/3319535.3354212

[32] Ron Wakkary, Audrey Desjardins, Sabrina Hauser, and Leah Maestri. 2013. A Sustainable Design Fiction: Green Practices. *ACM Trans. Comput.-Hum. Interact.* 20, 4, Article 23 (Sept. 2013), 34 pages. https://doi.org/10.1145/2494265

[33] Sebastian Zimmeck and Steven M. Bellovin. 2014. Privee: An Architecture for Automatically Analyzing Web Privacy Policies. In *23rd USENIX Security Symposium (USENIX Security 14)*. USENIX Association, San Diego, CA, 1–16. https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/zimmeck

[34] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel Reidenberg, N Cameron Russell, and Norman Sadeh. 2019. MAPS: Scaling privacy compliance analysis to a million apps. *Proceedings on Privacy Enhancing Technologies* 2019, 3 (2019), 66–86.

[35] Yixin Zou, Shawn Danino, Kaiwen Sun, and Florian Schaub. 2019. You 'Might' Be Affected: An Empirical Analysis of Readability and Usability Issues in Data Breach Notifications. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3290605.3300424

Received May 2021; revised September 2021; accepted October 2021