9

Contextual Integrity as a Gauge for Governing Knowledge Commons

Yan Shvartzshnaider,[1] *Madelyn Rose Sanfilippo,*[2] *and Noah Apthorpe*[3]

9.1 INTRODUCTION

This chapter describes our approach to combine the Contextual Integrity (CI) and Governing Knowledge Commons (GKC) frameworks in order to gauge privacy expectations as governance. This GKC-CI approach helps us understand how and why different individuals and communities perceive and respond to information flows in very different ways. Using GKC-CI to understand consumers' (sometimes incongruent) privacy expectations also provides deeper insights into the driving factors behind privacy norm evolution.

The CI framework (Nissenbaum, 2009) structures reasoning about the privacy implications of information flows. The appropriateness of information flows is defined in context, with respect to established norms in terms of their values and functions. Recent research has operationalized CI to capture users' expectations in varied contexts (Apthorpe et al., 2018; Shvartzshnaider et al., 2016), as well to analyze regulation (Selbst, 2013), establish research ethics guidelines (Zimmer, 2018), and conceptualize privacy within commons governance arrangements (Sanfilippo, Frischmann, and Strandburg, 2018).

The GKC framework examines patterns of interactions around knowledge resources within particular settings, labeled as action arenas, by identifying background contexts; resources, actors, and objectives as attributes; aspects of governance; and patterns and outcomes (Frischmann, Madison, and Strandburg, 2014). Governance is further analyzed by identifying strategies, norms, and rules-in-use through an institutional grammar (Crawford and Ostrom, 1995). According to GKC,

[1] Assistant Professor/Faculty Fellow in the Courant Institute of Mathematical Sciences, NYU; Visiting Associate Research Scholar at the Center for Information Technology Policy (CITP), Princeton University.

[2] Assistant Professor, School of Information Sciences, University of Illinois at Urbana-Champaign; Affiliate Scholar, The Vincent and Elinor Ostrom Workshop in Political Theory and Policy Analysis, Indiana University, Bloomington. Ph.D., Indiana University, Bloomington; M.I.S., Indiana University, Bloomington; B.S., University of Wisconsin-Madison.

[3] Assistant Professor, Department of Computer Science, Colgate University; Ph.D., Department of Computer Science, Princeton University; Graduate Student Fellow, Center for Information Technology Policy, Princeton University.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

strategies are defined in terms of attributes, aims, and conditions; norms build on strategies through the incorporation of modal language; and rules provide further structure by embedding norms with consequences to sanction non-compliance. For example, a strategy can describe a digital personal assistant that uses audio recordings of users (attributes) in order to provide personalized advertisements (aim) when a user does not pay for an ad-free subscription (condition). If this information flow also included modal language, such as a hedge, like "may" and "could," or a deontic, like "will" and "cannot," it would be a norm. The addition of a consequence, such as a denial of service or financial cost, would make this example a rule. It is also notable that, from this perspective, there are differences between rules-on-the-books, which prescribe, and rules-in-use, which are applied.

GKC and CI are complementary frameworks for understanding privacy as both governing institutions (Sanfilippo, Frischmann, and Strandburg, 2018) and appropriate flows of personal information, respectively. Within the GKC framework, as with the broader intellectual tradition of institutional analysis, an institutional grammar can be applied to deconstruct individual institutions (Crawford and Ostrom, 1995). Table 9.1 illustrates the overlap between these frameworks and how each provides parameter specificity to the other. While the CI framework deconstructs information flows, the GKC framework considers governance structures and constraints regarding actors and their interactions with knowledge resources. Consider the digital personal assistant example from the previous paragraph. Under the institutional grammar (Crawford and Ostrom, 1995), the digital personal assistant, audio recordings, and users are all considered "attributes." The CI framework further divides these elements into sender, information type and subject parameters, respectively. Conversely, the CI framework uses the "transmission principle" parameter to articulate all constraints on information flows, while the GKC framework provides definitions of aims, conditions, modalities, and consequences.

In this work, we use the GKC and CI frameworks to understand the key aspects behind privacy norm formation and evolution. Specifically, we investigate divergences between privacy expectations and technological reality in the IoT domain.

TABLE 9.1 *Conceptual overlap between CI and Institutional Grammar (GKC) parameters*

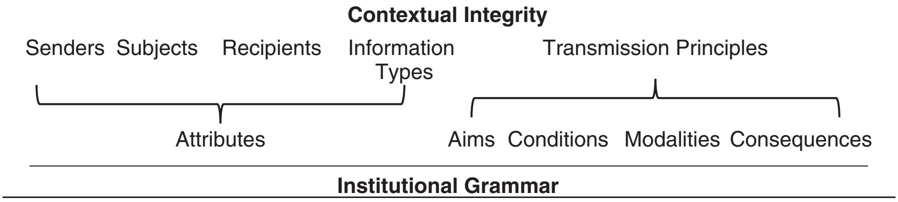

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

The consumer Internet of things (IoT) adds Internet-connectivity to familiar devices, such as toasters and televisions, resulting in data flows that do not align with existing user expectations about these products. This is further exacerbated by the introduction of new types of devices, such as digital personal assistants, for which relevant norms are only just emerging. We are still figuring out whether the technological practices enabled by these new devices align with or impact our social values. Studying techno-social change in the IoT context involves measuring what people expect of IoT device information flows as well as how these expectations and underlying social norms emerge and change. We want to design and govern technology in ways that adhere to people's expectations of privacy and other important ethical considerations. To do so effectively, we need to understand how techno-social changes in the environment (context) can lead to subtle shifts in information flows. CI is a useful framework for identifying and evaluating such shifts as a gauge for GKC.

We conduct a multi-part survey to investigate the contextual integrity and governance of IoT devices that combines open-ended and structured questions about norm origins, expectations, and participatory social processes with Likert-scale vignette questions (Apthorpe et al., 2018). We then perform a comparative analysis of the results to explore how variations in GKC-CI parameters affect privacy strategies and expectations and to gauge the landscape of governing norms.

9.2 RESEARCH DESIGN

In the first part of the survey, we asked respondents to list the devices they own and how they learn about the privacy properties of these devices (e.g., privacy policies, discussions with legal experts, online forums). We next presented the respondents with scenarios A through D, as described in Table 9.2, each scenario was followed by applied questions based on the GKC framework.

Each scenario focused on different factors that previous research has identified as having an effect on users' expectations and preferences (Apthorpe et al., 2018). Scenario A focused on third-party information sharing practices involving a smart TV that tracks viewing patterns and TV watching habits that are sold to an advertiser. Questions assessed the respondents' specific concerns in this scenario as well as their anticipated reactions. We interpreted these reactions as indicators of respondents' privacy expectations and beliefs as well as their understanding of information flows in context.

The remaining scenarios were built on Scenario A to explore different factors affecting privacy opinions and reactions. Scenario B introduced an additional, exogenous influence: a parallel, cross platform tracking incident that happened to someone else the respondent might know. Questions assessed how experiences with cross-device information flows and surrounding factors alter respondents' expectations and resulting actions. This provides a sense of communities and contexts surrounding use, in order to support future institutionalization of information flows to better align with users' values.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

TABLE 9.2 *Survey scenarios with corresponding aspects of the GKC framework*

#	Scenario	GKC Aspects
A	Imagine you're at home watching TV while using your phone to shop for socks on Amazon. Your TV then displays an ad informing you about a great discount on socks at a Walmart close to your neighborhood.	**Background:** normative values **Attributes:** resources **Patterns and Outcomes:** benefits
B	You later hear from your neighbor that a similar thing happened to him. In his case, his wife posted on Facebook about their dream vacation. A few days later he noticed an ad as he was browsing the web from a local cruiser company.	**Background:** normative values **Attributes:** resources, community members, goals and objectives **Governance:** institutions **Patterns and Outcomes:** benefits
C	Companies usually detail their information handling practices in their privacy policies and terms of service. Imagine you do read through the privacy policy for your smart TV. You find a statement saying that the TV could, sometimes, send your information to third parties for analysis to offer you all the top features. The privacy policy also states that you may disable third party sharing; however, this may cause additional subscription charges for some features.	**Governance:** context, institutions, actors **Patterns and Outcomes:** benefits, costs, legitimacy
D	You have an acquaintance who is a software engineer. They tell you that you shouldn't be concerned. It's considered a normal practice for companies to track the habits and activities of their users. This information is then typically sold to third parties. This is how you can get all of these free personalized services!	**Attributes:** community members, goals and objectives **Governance:** institutions, actors **Patterns and Outcomes:** costs, legitimacy

Scenario C focused on privacy policies and whether they mitigate privacy concerns. Specifically, we asked how often respondents read privacy policies and what they learn from them. We also queried whether the practice of information sharing with third parties potentially changes respondents' behavior whether or not the data are anonymized. Finally, we asked whether the respondents would be willing to employ a workaround or disable information sharing for an additional charge – examples of rules-in-use contrasting sharply with rules-on-the-books that otherwise support information flows respondents may deem inappropriate.

Scenario D assessed how exogenous decision-makers influence privacy perceptions and subsequent behavior by providing respondents with an example of expert advice. Questions about this scenario addressed differences in perceptions between stakeholder

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

TABLE 9.3 *Smart home GKC-CI parameters selected for information flow survey questions*

Sender	Modality	Aim
Google Home	can	if the information is used for advertising
Amazon Echo (Alexa)	might	if the information is used for academic research
Apple HomePod (Siri)	will	if the information is used for developing new device features
Smart watch		
Garmin watch		
Subject & Type		Condition
Your personal information		if you have given consent
		if you are notified
Your location		
Recorded audio		
Recipient		Consequence
Its manufacturer		if the information is used to generate summary statistics
A third party		if the information is necessary for the device to function properly
		if the information is used to personalize content

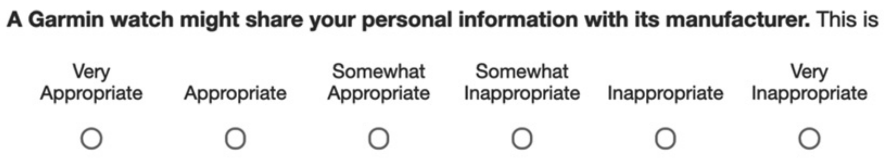

A Garmin watch might share your personal information with its manufacturer. This is

Very Appropriate	Appropriate	Somewhat Appropriate	Somewhat Inappropriate	Inappropriate	Very Inappropriate
○	○	○	○	○	○

FIGURE 9.1 Example baseline information flow question

groups as well as the legitimacy of expert actors in governance. While Scenario D specifically included a software engineer as the exemplar expert role, a parallel study has assessed perceptions of many additional expert actors (Shvartzshnaider, Sanfilippo, and Apthorpe, under review).

The second section of the survey tested how variations in CI and GKC parameters affect the perceived appropriateness of information flows. We designed this section by combining GKC parameters with an existing CI-based survey method for measuring privacy norms (Apthorpe, 2018).

We first selected GKC-CI parameters relevant to smart home device information collection. These parameters are listed in Table 9.3 and include a variety of timely privacy issues and real device practices.

The questions in this section followed a parallel structure. Respondents were first presented with an information flow description containing a randomly selected combination of sender, subject, information type, recipient, and modal parameters (Figure 9.1). Respondents rated the appropriateness of this flow on a 6-point Likert scale from "very inappropriate" to "very appropriate."

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

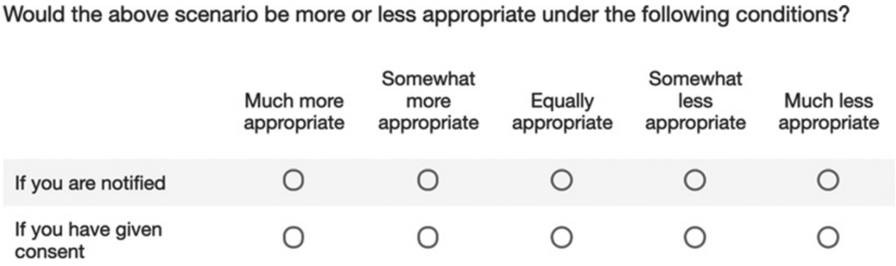

Would the above scenario be more or less appropriate under the following conditions?

	Much more appropriate	Somewhat more appropriate	Equally appropriate	Somewhat less appropriate	Much less appropriate
If you are notified	○	○	○	○	○
If you have given consent	○	○	○	○	○

FIGURE 9.2 Example question with varying condition parameters

This baseline question was followed by a series of matrix-style multiple choice questions with one row for each condition, aim, and consequence parameter (Figure 9.2). Respondents were asked to indicate how each of these parameters would affect the appropriateness of the original information flow on a 5-point Likert scale from “much more appropriate” to “much less appropriate.”

This process was repeated three times for each survey participant. Each participant rated three sets of baseline flows with different subject/type/recipient/modal parameters and corresponding matrices for condition/aim/consequence parameters. Null parameters were included as controls for each category.

The survey concluded with a series of standard demographics questions, querying respondents’ age, gender, state of residence, education level, and English proficiency. Each of these questions had a “prefer not to disclose” option in case respondents were uncomfortable divulging this information.

We created the survey using Qualtrics. We conducted “cognitive interviews” to test survey before deployment via UserBob, an online usability testing platform. Five UserBob workers were asked to take the survey while recording their screen and providing audio feedback on their thought processes. These workers were paid $1 per minute, and all completed the survey in less than 10 minutes. While the UserBob responses were not included in the results analysis, they confirmed the expected survey length of less than 10 minutes and that the survey did not contain any issues that would inhibit respondents’ understanding.

We deployed the survey as a Human Intelligence Task (HIT) on Amazon Mechanical Turk (AMT). The HIT was limited to AMT workers in the United States with a 90–100 percent HIT approval rating. We recruited 300 respondents and paid each $1 for completing the survey.

We began with 300 responses. We then removed 14 responses from individuals who provided incomprehensible answers or non-answers to the free-response questions. We also removed 2 responses from individuals who answered all matrix questions in the same column. This resulted in 284 total responses for analysis.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

9.3 GOVERNING APPROPRIATE FLOWS OF PERSONAL INFORMATION FROM CONSUMER IOT

We analyze our survey results from the combined GKC-CI perspective. We use GKC framework to identify the background environment (specific context) of consumer IoT, attributes involved in the action arena of IoT information flows (including goals and objectives), governance rules within consumer IoT contexts, and various patterns and outcomes, including the perceived cost and benefits of IoT information flows. We also use the CI framework with the institutional grammar parameters (aims, conditions, consequences, modalities) as transmission principles to understand what specific aspects of governance have the most significant impact on respondent perceptions.

9.3.1 *Background Environment*

Smart devices are pervasive in Americans' lives and homes. We interact with a wide range of these supposedly smart systems all the time, whether we recognize and consent to them or not, from Automated License Plate Readers (ALPR) technologies tracking drivers' locations (Joh, 2016) to Disney World's MagicBand system (Borkowski et al., 2016) to Alexa in dorm rooms (Manikonda et al., 2018). These devices, which are part of a larger digital networked environment, collect massive amounts of data that surreptitiously capture human behaviors and support overt sociotechnical interactions in public and private spaces.

It is notable that there are very different scales of use and applications of smart devices, with many deployed publicly without public input. In contrast, smart devices in individuals' homes are most often configured by the users themselves with appropriate use negotiated within households. Notable exceptions include the controversial and well-publicized implementations of smart locks and systems in rental housing (e.g., Geeng and Roesner, 2019) and uses of IoT to surveil victims by perpetrators of domestic violence (Tanczer et al., 2018). These consumer IoT devices have wildly different patterns of interactions and governance. They are operated under complex arrangements of legal obligations, cultural conditions, and social norms without clear insight into how to apply these formal and informal constraints.

It is thus important to establish applicable norms and evaluate rules-in-use to support good governance of consumer IoT moving forward. Understanding interactions where users have some control of institutional arrangements involving their devices is helpful toward this end. We therefore focus on consumers' everyday use of smart devices, primarily smartphones, wearables, and in-home smart devices. It is our objective to understand both how users would like information flows associated with these devices to be governed and how their privacy perceptions are formed.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

The background context for personal and in-home IoT device use extends beyond individual interactions with smart devices. It includes aggregation of information flows from devices and interactions between them, discussion about the relevant normative values surrounding device use, and governance of information flows. There are distinct challenges in establishing norms, given that there is no default governance for data generated, as knowledge resources, or predictable patterns of information to help form user expectations.

Our survey respondents documented the devices they owned, which aligned with recent consumer surveys of IoT prevalence (e.g., Kumar et al., 2019). About 75 percent of respondents reported owning more than one smart device, with 64 percent owning a smart TV and 55 percent owning a Digital Personal Assistant (such as an Amazon Echo, Google Home, or Apple HomePod). Smartwatches were also very popular. A small percentage of respondents owned smart lightbulbs or other Internet-connected appliances.

As these devices become increasingly popular and interconnected, the contexts in which they are used are increasingly complex and fraught with value tensions, making it important to further study user preferences in order to develop appropriate governance. For example, digital personal assistants don't clearly replace any previous analogous devices or systems. They therefore lack pre-existing norms or underlying values about appropriateness to guide use. In contrast, smart televisions are obviously analogous to traditional televisions and are thereby used in ways largely guided by existing norms. These existing norms have often been shaped by entrenched values but do not always apply to emerging information flows from and to new smart features. The resulting tensions can be resolved by identifying relevant values and establishing appropriate governing institutions around IoT information flows. To do so, it is first important to understand the relevant factors (attributes) so as to clarify how, when, and for what purpose changes in information flows governance are and are not appropriate.

9.3.2 *Attributes*

9.3.2.1 Resources

Resources in the IoT context include both (1) the data generated by devices and (2) knowledge about information flows and governance. The latter also includes characteristics of these devices, including necessary supporting technologies and personal information relevant to the IoT action arena.

The modern home includes a range of devices and appliances with Internet-connectivity. Some of these devices are Internet-connected versions of existing appliances, for example, refrigerators, TVs, thermostats, lightbulbs. Other devices, such as digital assistants (e.g., Amazon Echo and Google Home), are new. These devices produce knowledge by generating and consuming information flows. For

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

example, a smart thermostat uses environmental sensors to collect information about home temperature and communicates this information to cloud servers for remote control and monitoring functionality. Similar information flows across devices are causing the IoT ecosystem to evolve beyond the established social norms. For example, now refrigerators order food, toasters tweet, and personal health monitors detect sleeping and exercise routines. This rapid change in the extent and content of information flows about in-home activities leads to a mismatch between users' expectations and the IoT status quo. Furthermore, mismatches extend beyond privacy to features, as some new "smart" functions are introduced for novelty sake, rather than consumer preferences, such as kitchen appliances that are connected to social media.

Our survey respondents' comments reveal discrepancies between users' privacy perceptions/preferences and how IoT devices are actually used. This provides further insight into the attributes of data resources within this context by illustrating what is considered to be appropriate. For example, some respondents noted that even though they have smart TVs, they disconnect them from the Internet to limit communication between devices. Generally, our survey results highlight the range of confusion about how smart devices work and what information flows they send.

A few respondents implied that they were only learning about IoT cross-device communications through the scenarios described in our survey, describing their surprise (e.g., "How they already know that. How did it get from my phone to the tv? That seems very fishy") or in some cases absolute disbelief ("I see no connection between what I'm doing on the phone and a random TV ad") that such a thing was possible. One respondent specifically summarized this confusion amidst common experiences with new technologies:

> *At first, you are concerned. The lightning fast speed at which Google hits you in the heads* [sic] *for an item you were considering buying makes you believe they are spying on you. They aren't spying, because spying implies watching you without your permission, but in using the service you give them complete permission to use any data you put into search queries, posts, etc, to connect you to items you are shopping for, even if it is just to look around.*
>
> *Social media consumers do not understand that they are NOT the customer. They are the product. The customer is the numerous businesses that pay the platform (Google, Facebook, etc) various rates to get their product in front of customers most likely to pay. Radio did this long before Cable TV who did this long before Social Media companies. It's a practice as old as steam.*

This quotation highlights perceived deception about information collection practices by popular online platforms and IoT devices. Users of IoT devices are shaping their expectations and practices amidst a lack of transparency about privacy and problematic notions of informed consent (e.g., Okoyomon et al., 2019). This respondent also touches on the inextricable links between the two knowledge

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

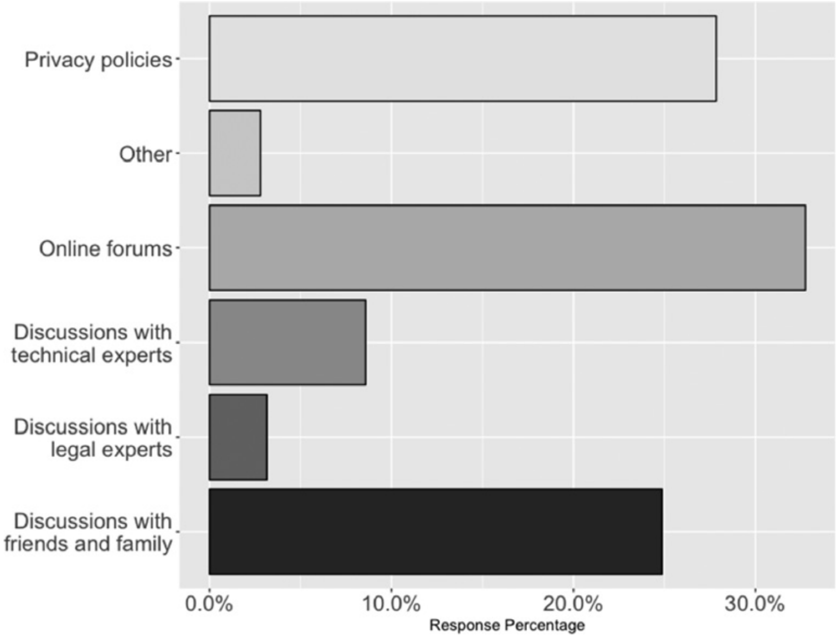

FIGURE 9.3 Where respondents learn about the privacy implications of IoT devices

resources; when users have poor, confusing, or limited explanations of information flows, they fail to understand that they are a resource and that their data is a product.

As Figure 9.3 illustrates, respondents learn about IoT information flows and privacy from a variety of different sources. Online forums represent the most prevalent source of privacy information, yet only just over 30 percent of respondents turn to online forums of IoT users with privacy questions. Privacy policies and discussions with friends and family were also common sources of privacy information, but even these were only consulted by 28 percent and 25 percent of respondents, respectively. Respondents turned to technical and legal experts for privacy information even less frequently, with only 9 percent and 3 percent of respondents reporting these sources, respectively. Overall, there was no single source of privacy information consulted by a majority of respondents.

9.3.2.2 Community Members

Community members, through the lens of the GKC framework, include those who participate and have roles within the action arena, often as users, contributors, participants, and decision-makers. The action arena also includes a variety of additional actors who shape these participants' and users' expectations and preferences, including lawyers and privacy scholars; technologists, including engineers

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

and developers; corporate social media campaigns; anonymous discussants in online forums; and friends and family, which we examine in a related study (Shvartzshnaider, Sanfilippo, and Apthorpe, under review). It is important to consider who is impacted, who has a say in governance, and how the general public is impacted. In this context, community members include IoT device owners, developers, and users, as well as users' family, friends, and neighbors in an increasingly connected world.

While the respondents who depend on online communities and forums for privacy information are a small subset, those communities represent an important source of IoT governance in use. User-generated workarounds and privacy discussions are meaningful for understanding and establishing appropriate information flows. Users are thus the community-of-interest in this context, and those who responded to our survey reflect the diversity of users. The respondents were 62 percent male and 37 percent female with an average age of 34.5 years. 53 percent of the respondents had a Bachelor's degree or higher. 38 percent of respondents self-reported annual incomes of <$40,000, 43 percent reported incomes of <$80,000, 8 percent reported incomes of <$100,000, and 10 percent reported income of > $100,000. We have not established clear demographic indicators for the overall community of IoT users, in this sense, beyond ownership and a skew toward a younger population. However, it is also possible that tech savviness is overrepresented among users.

9.3.2.3 Goals and Objectives

Goals and objectives, associated with particular stakeholders, are grounded in history, context, and values. It is important to identify the specific obstacles and challenges that governance seeks to overcome, as well as the underlying values it seeks to institutionalize.

In our studies, the respondents identified multiple governance objectives and dilemmas associated with information flows to and from IoT devices, including control over data collection and use, third parties, and autonomy in decision-making. Interests among respondents were split between those who valued cross-device information flows and those who felt increased interoperability and/or communication between devices was problematic. Additionally, there were a few respondents who agreed with some of the perceived interests of device manufacturers that value monetization of user data; these respondents appreciated their ability to utilize "free services" in exchange for behavioral data collection. Furthermore, there are additional tensions between the objectives of manufacturers and developers and the interests of users, as evidenced by the split in trust in the expertise of a technical expert in judging appropriateness of information flows. These results show fragmentation in perception of both governance and acceptance of the status quo for information flows around IoT devices.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

9.3.3 *Governance*

Through the lens of the GKC framework, including the institutional grammar, we gain insight into different aspects of governance. We can capture how the main decision-making actors, individual institutions, and the norms governing individual information flows emerge and change over time, as well as how these norms might be enforced. Results also indicate that privacy, as appropriate flows of personal information, governs interactions with and uses of IoT devices. For example, we see evidence that anonymization, as a condition modifying the information type and its association with a specific subject within an information flow, does not serve as meaningful governance from the perspective of respondents. Fifty-five percent of respondents stated that they would not change their behavior, or support cross-device communication, just because data was anonymized. It is not immediately clear, from responses to that question alone, what leads to divergence on this interpretation of anonymization or any other perceptions about specific information flows. However, it echoes theorization about CI that incomplete transmission principles are not helpful in understanding information flows (e.g., Bhatia and Breaux, 2018), extending this idea to governance; the condition of anonymity is not a stand-alone transmission principle.

This aligns with our approach combining the GKC and CI frameworks to gauge the explicit and implicit norms that govern information flows within a given context. The CI framework captures norms using five essential parameters of information flows. Four of the parameters capture the actors and information type involved in an information flow. The fifth parameter, transmission principle, constrains information flows. The transmission principle serves as a bridging link between the CI and GKC frameworks. Figure 9.4 shows the average score for perceived appropriateness for an information flow without qualifying it with the transmission principle. We remind the reader that the respondents were first presented with information flow descriptions using sender, subject, information type, recipient, and modal parameters. They rated the appropriateness of these flows on a 6-point Likert scale from "very inappropriate" (-2) to "very appropriate" (+2).

For the GKC framework questions in the first part of the survey, 73 percent of respondents reported that they would change their behaviors in response to third-party sharing. Specific actions they would take are illustrated in Figure 9.6. Figure 9.4 shows that respondents view a "manufacturer" recipient less negatively than a generic third party. Additionally, not stating a recipient all together has a lesser negative effect on information flow acceptability than a generic "third party" recipient. We can speculate that when the recipient is omitted, the respondents mentally substitute a recipient that fits their internal privacy model, as shown in previous research (Martin and Nissenbaum, 2016).

We further gauge the effect on user perceptions of aims, conditions, modalities, and consequences as components of transmission principles. Figure 9.5 illustrates

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

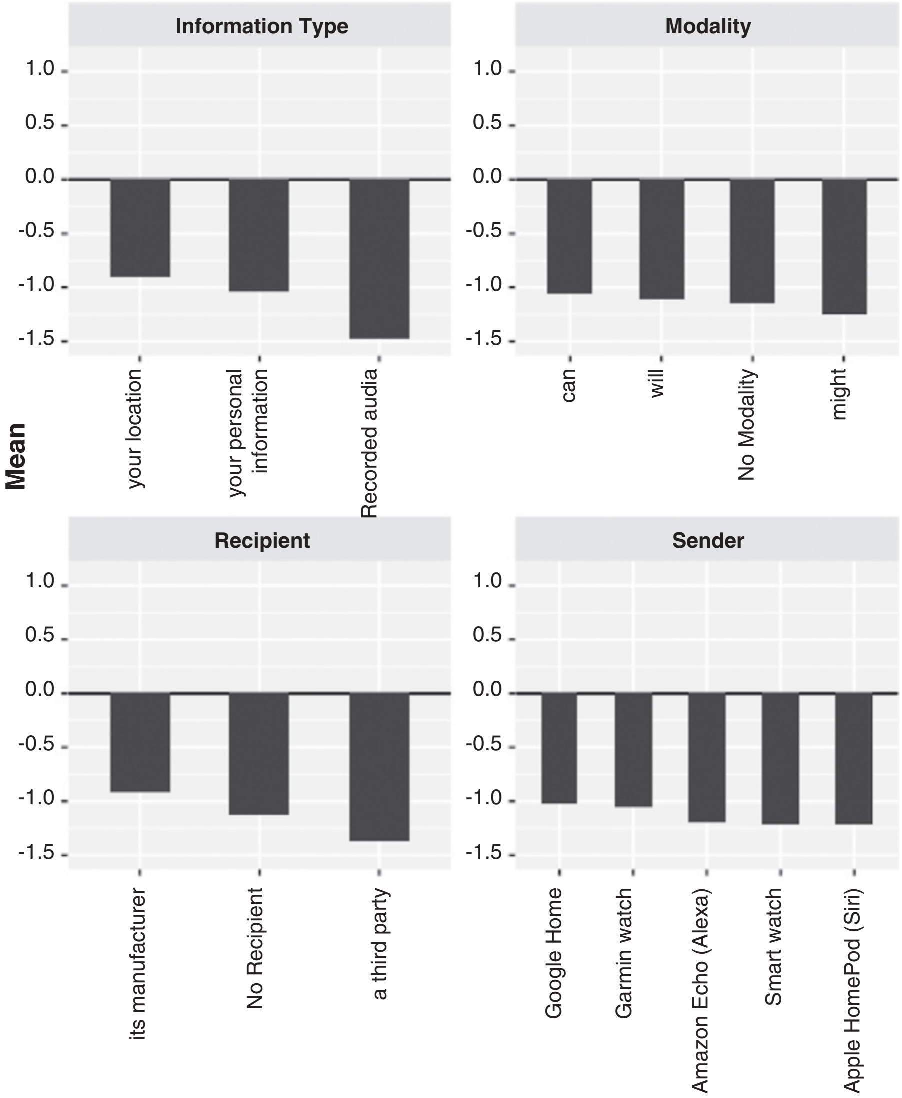

FIGURE 9.4 Average perceptions of information flows by parameter
This figure illustrates the average participant opinion of information flows controlled to specific examples of information type and subject, modalities, recipients, and senders.

changes in average perceptions based on the addition of specific aims, conditions, and consequences to the description of an information flow. We see that stating a condition (such as asking for consent, upon notification or keeping the data anonymous) has a positive effect on the perception of appropriateness. Conversely, we see that not stating an aim correlates with positive perception,

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

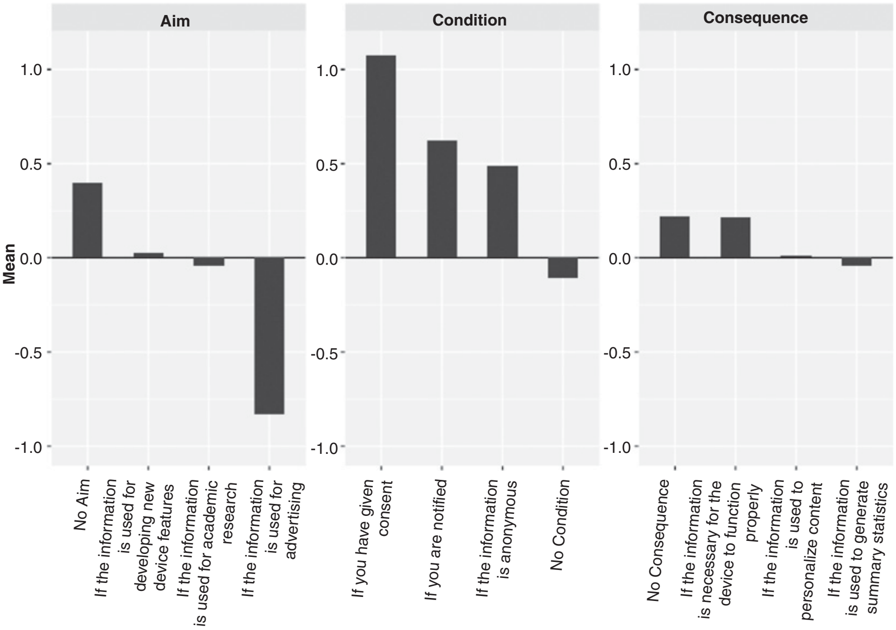

FIGURE 9.5 The impact of specific parameters in changing respondent perceptions of information flows. This figure indicates the average change in perceptions in response to specific examples for each parameter. It does not indicate initial perceptions, in contrast to Figure 9.4.

while the respondents seemed on average neutral towards "for developing new features" and "for academic research" aims, they show negative attitude towards the "for advertising purposes" aim. When it comes to consequences, the results show that the respondents view not stating a consequence as equal, on average, to when the information "is necessary for the device to function properly." However, respondents viewed information flows with the consequence "to personalize content" slightly positively, while viewing information flows with the consequence of "[generating] summary statistics" correlates with slightly negative perception.

Respondents also identified a number of additional approaches that they might take in order to better control flows of their personal information and details of their behaviors between devices. In addition to browsing anonymously and disconnecting their smart TV from the Internet, various respondents suggested:

- *"Use a VPN"*
- *"Wouldn't buy the TV in the first place"*
- *"It's just getting worse and worse. I'll almost certainly return it."*
- *"Research and see if there is a way to not have my info collected."*

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

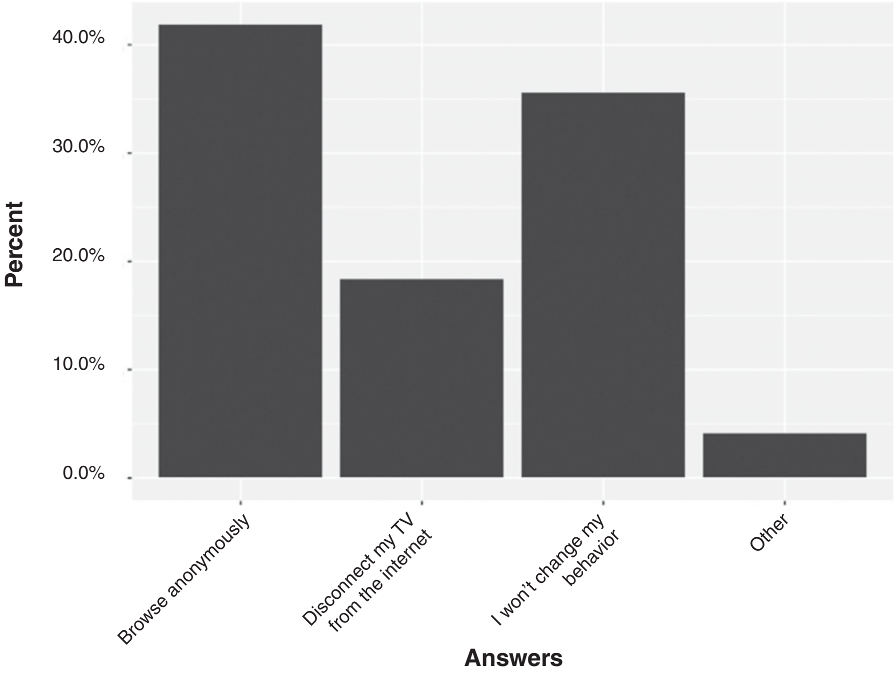

FIGURE 9.6 User actions in response to third-party sharing scenarios

- *"Be much more careful about my browsing/viewing habits."*
- *"Circumvent the tracking"*
- *"Try to find a way to circumvent it without paying"*
- *"Sell it and get a plain TV"*
- *"Block access to my information"*
- *"Delete cookies"*
- *"Disable features"*

When they perceived information flows to be inappropriate, many respondents favored rules-in-use that would circumvent inadequate exogenous governance. While many respondents favored opportunities to opt out of inappropriate flows, a small sub-population developed their own approaches to enact their privacy preferences as additional layers of governance in use. Often these work-arounds subverted or obfuscated default information flows.

9.3.3.1 Rules-in-Use and Privacy Policies

Few respondents found the rules-on-books described in privacy policies to be useful for understanding information flows associated with IoT devices. Many respondents

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

described how they found privacy policies lengthy and confusing. For example, when asked what they learn from reading privacy policies, one respondent explained:

> *That they* [sic] *hard to read! Seriously though, they are tough to interpret. I know they try and protect some of my information, but also share a bunch. If I want to use their services, I have to live that* [sic].

One of the 62 respondents who reported that they read privacy policies "always" or "most of the time" further elaborated:

> *I've learned from privacy policies that a lot of times these company [sic] are taking possession of the data they collect from our habits. They have the rights to use the information as they pleased, assuming the service we're using from them is advertised as 'free' I've learned that sometimes they reserve the right to call it their property now because we had agreed to use their service in exchange for the various data they collect.*

The information users learn from reading a privacy policy can undermine their trust in the governance imposed by IoT device manufacturers. The above comment also touches on issues of data ownership and rights to impact or control information flows. Privacy policies define rules-on-the-books about these issues, which some respondents perceive to be imposed governance. However, as noted by another respondent, the policies do not apply consistently to all devices or device manufacturers:

> *That companies can be pretty loose with data; that some sell data; that others don't go into specifics about how your data is protected; and there are some that genuinely seem to care about privacy.*

This comment emphasizes an important point. For some respondents, practices prescribed in privacy policies affect how they perceive each respective company. In cases where privacy policy governance of information flows aligns with social norms, companies are perceived to care about privacy. Respondents also identify these companies as more trustworthy. In contrast, privacy policies that are vague about information flows or describe information flows that respondents perceive to be egregious or excessive, such as selling user data to many third parties, indicate to respondents that associated companies do not care about user privacy.

Relative to these inappropriate and non-user centered information flows and policies, respondents also described rules-in-use and work-arounds that emerged in order to compensate for undesirable rules-on-the-books. Over 80 percent of respondents indicated that they would pursue work-arounds, with many pursuing alternate strategies even if it took an hour to configure (31 percent) or regardless of difficulty (32 percent).

A few respondents recognized that privacy policies sometimes offer ways to minimize or evade tracking, such as outlining opportunities to opt out, as well as

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

defining the consequences of those choices. When asked "What do you learn from privacy policies?," one respondent elaborated:

> *Occasionally, there are ways to minimize tracking. Some of the ways the data is used. What things are needed for an app or device.*

In this sense, privacy policies disclose and justify information flows, often discouraging users from opting-out through institutionalized mechanisms, such as options to disable recommendations or location services, by highlighting the features they enable or the consequences of being left out. However, despite institutionalized mechanisms to evade tracking, opt out options are sometimes insufficient to protect privacy (Martin, 2012). Furthermore, many respondents don't actually read privacy policies and therefore may not be aware of them. Thus, individuals also develop their own approaches and share them informally among friends and online communities, as shown in Figure 9.1.

Through the lens of the GKC framework, privacy policies serve as a source for rules-on-the-books. These rules govern the flow of information into and out of IoT companies. From respondents' comments, we see that privacy policies play an important role in shaping their expectations for better for worse. On one side, the respondents turn to privacy policies because they want to learn "what [companies] do and how they may use information they receive." On the other side, respondents echoed the general public frustration of not being able to "to learn anything because [privacy policies] are purposefully wordy and difficult to understand." Companies that outline clear information governance policy help inform users' expectations about their practices, while those companies that offer ambiguous, lengthy, hard to understand policies force users to rely on their existing (mostly negative) perceptions of company practices and/or turn to other sources (family, experts) for information.

Finally, the respondents discuss several options for dealing with the gap between rules-on-the-books and their expectations. First, they could adjust their expectations ("these smart devices know too much about me," "be more careful about what I do searches on"). They could also find legal ways to disable practices that do not align with their expectations, such as paying to remove ads or changing settings ("I trust them but I still don't like it and want to disable"). In addition, they could opt out from the service completely ("Sell it and get a plain TV").

9.3.4 *Patterns and Outcomes*

Our survey reveals a significant fragmentation within the community of IoT users relative to current governance practices, indicating irresolution in the action arena. As we piece together data on who IoT users are and how they are shaping appropriate flows of personal information from and between their smart devices, certain patterns and outcomes become evident. Table 9.4 illustrates how respondents' preferences about third party sharing, professed concerns about privacy, and device ownership

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

TABLE 9.4 *Average perceptions of information flow appropriateness gauged by respondent subpopulations. For each subcommunity we calculate the number of respondents and the average perception score across information flows including consequence, condition, and aim.*

	Embrace Tech (own >2 devices)	Don't embrace tech	Concerned about third party sharing	Not concerned about third party sharing
Unconcerned	0.53 (n=48)	0.5 (n=35)	0.5 (n=52)	0.6 (n=31)
Concerned	0.06 (n=94)	0.05 (n=92)	0.05 (n=171)	0.7 (n=15)

shape their average perceptions of governance outcomes around IoT. We assessed the extent to which respondents embraced technology based on the number of devices they own.

Table 9.4 divides the respondents of our survey into subcommunites based on the opinions of various IoT practices elicited from the first part of the survey. Some respondents largely have embraced IoT technology[4] and are not concerned about privacy issues.[5] Others, while embracing the technology, are concerned about privacy issues. Concerns about third party sharing or a lack of embrace of smart technology yield very different opinions, on average. We cluster these subcommunities into three groups, in order to gauge their perceptions.

When gauging the respondents' perceptions, we note that those who are unconcerned about the privacy implications of cross platform sharing, regardless of other values associated with information governance, have on average favorable views of information flows. Additionally, those respondents who express general concern about the privacy implications, but are not concerned about third party sharing, have similar perceptions on average. These subpopulations of our respondents are the most likely to belong to group 1, who perceive current governance of IoT information flows to be positive, on average. In contrast, those who are concerned about privacy and either don't embrace smart technologies or express concerns about third party sharing are most likely to belong to group 3, who are slightly dissatisfied with current governance outcomes on average. Finally, group 2 is generally concerned about privacy but embraces smart devices with average perceptions slightly above neutral.

We now highlight the open-ended comments from respondents belonging to each group, that put their opinions in context, in an effort to better understand fragmentation and what underlying beliefs and preferences lead to divergent normative

[4] Respondents indicated to own more than two smart devices.

[5] Respondents in Group 1 indicated that they weren't concerned with the privacy implications of the survey Scenario A.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

patterns. While individual comments are not representative, they illuminate individuals' rationales underlying perceptions associated with groups.[6]

9.3.4.1 Group 1: Positive Perceptions

This group includes respondents that positively perceive information sharing practices and tend to emphasize both user consent and preferences for personalization on average. As one user specified:

> *Because I knew what I was getting myself into when using these types of products. How else should companies market to me? We could go back to the old days when there was no personalization at all, when ads were completely dumb and never actually spoke to your needs. But, how is that better? People worry about privacy, but they should only really be concerned with security, which is not the same thing. Keep my financial info secure, keep certain embarrassing stuff under wraps to the public, sure, but we share so much of our lives openly that it seems silly to scoff at ad personalization. I do, however, get annoyed when it doesn't seem personalized ENOUGH, because then it's akin to the uncanny valley for CGI … in those moments, I'm frustrated that the personalization isn't stronger, such as when I continually see ads for stuff I've already bought.*

Some participants in this group also expressed a firm belief that linking devices that share data would have to be deliberate on the part of users. These users would implicitly consent to information flows, in contrast to respondents with neutral and negative perceptions. In this sense, discoverability, or the ability of smart devices to recognize and communicate with one another, was not acknowledged as a smart feature. For example:

> *For the devices to work like that I must have linked them in some way. That deliberate action would have been my consent to allow these devices to exchange data.*

9.3.4.2 Group 2: Neutral Perceptions

Respondents in this group have a relatively neutral perception of information flows on average. While participants in this group seem to recognize the issues related to discoverability between devices, they don't view them as a privacy violation. As one participant explained their thought process:

> *I feel like at this point everything is somehow connected. There have been many times where I browse the internet and then on my Facebook profile I see adds for the very*

[6] Each group was identified by their average perceptions of appropriateness, rather than by similarity in open-ended responses.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

> *thing that I was looking for. I know that it is an effort to target me and things that I might like, I don't think my privacy is being compromised.*

They accept data flows between devices, relative to their user behavior, as standard practice and seem to perceive personalized advertising as independent of their privacy values. However, other members of this group raised concerns about the risks of specific information recipients:

> *I trust them because I think they just want to advertise to me better, I'd only be concerned if the information was being sold to criminals or hackers.*

In this sense, those with neutral perceptions of IoT information flows view credible commercial entities to be legitimate recipients. Sales and advertising are valid objectives, which various individuals within this moderate group saw as compatible with their privacy interests. In contrast, "criminals or hackers" were not seen to be acceptable recipients; future research should assess the differences in perceptions between these recipients and others.

In addition to concerns about some lesser-known third-party recipients, the past history of particular major manufacturers and commercial actors who have been careless or whose security has been compromised was also considered. Some respondents firmly believed that recent history with respect to breaches was unlikely to repeat, consistent with a recent study (Zou et al., 2018). One respondent explained their trust that violations of privacy would not recur:

> *because it seems that a lot of companies have gotten into trouble over the years and hopefully they're taking extra precautions these days.*

In other words, there is a belief that the companies would learn from past events and govern data in a way that was acceptable to them. This group was largely defined by acceptance of major manufacturers as trustworthy enough, without particular enthusiasm. Some of these users appeared to consider these flows in primarily the context of economic transactions.

9.3.4.3 Group 3: Negative Perceptions

Finally, those with negative perceptions of information flows and governance did not share the overall trust in companies to govern user data in accordance with social expectations. In particular, this group held negative perceptions of information flows between devices. Many of these respondents described these cross-platform flows as invasive:

> *It seems invasive and annoying. I also worry that my devices are exchanging information which each other that I didn't specifically agree to divulge. And who knows where else this information is going! All for what? To try and sell me garbage that I don't need and won't actually buy.*

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

The underlying problem was often with information being used out of context:

If it was just on the browser that I was using to search for socks, it wouldn't be as creepy. It's the fact that multiple platforms are being used in conjunction to analyze what I am doing for targeted advertising that I find invasive.

This sizeable community perceives current information flow practice and governance relative to IoT as violating their expectations.

Some respondents explained how IoT information flows also undermine their trust in other contexts because governance is non-transparent:

> *This seems like an invasion of privacy and makes me wonder what kinds of information it is collecting, storing, or otherwise utilizing for purposes not formally disclosed. Additionally, some devices are shared among families and friends when they visit. I find it to be a violation of my right to privacy to have data related to my phone searches and activities show up across multiple devices that are not used by only one person.*

This is only exacerbated by the industry's continued downplaying of the significance of data sharing.

This group of users was most unified and verbose in explaining their frustration with current governance and information flows in practice. They more often distrusted the technology industry and practitioners, such as in the software engineer scenario on our survey. In addition to not valuing personalization, some emphasized the problematic lack of control and uncertainty about data destinations beyond initial third-party recipients:

> *. . . who knows what happens to this data in the end? Will these third parties sell my info to other third parties? Of course they will. Is all this "free" stuff worth it? There's always a price, you know.*

Some respondents emphasized that current outcomes are egregious and that companies and regulators are falling short in governing user data:

> *I don't believe that it's something people should roll over about. When do we consider it out of hand? It's better to nip these kind of things in the bud. As a computer science major, having one persons opinion on the matter is not going to sway my opinion entirely. I wouldn't just get one opinion from a single doctor of my life was on the line would I?*

These respondents, in particular, emphasize that they want to play a more active role in governing their personal information flows.

Our results demonstrate the tensions that users experience when thinking of privacy in the IoT context. Through the scenarios addressing GKC concepts in the survey, we can observe divergence in interests and concerns of various respondents. Some welcome the new innovations and believe companies have their interest at heart. Others are more concerned, however, and often admit that they feel that

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

there is little they can do to protect their information. This reflects technological acceptance models in the larger population (e.g., Valdez and Ziefle, 2019). By gauging their perceived appropriateness of specific information flows, we can examine additional dimensions of governance using the language of the institutional grammar.

9.4 IMPLICATIONS

9.4.1 *Conceptual and Methodological*

As home environments evolve with the introduction of new technologies, norms of governance and information flow evolve as well. The growing tradition of GKC analysis of a cooperative governance schema offers a way to uncover the contributing elements related to a shift in privacy expectations.

Our approach relies on the GKC framework to identify emerging communities in a given context and then use the CI framework to pose questions about what information flows they consider appropriate. Our methodology bridges the two frameworks by quantifying the effect of each of the elements on the collective norms by measuring how each factor affects the appropriateness of information flows in a given context. This allows researchers to gauge the effect of various factors on the formation of the norms and could be employed to structure future case studies in other contexts to understand norm formation. Our study shows that omitting a condition has an effect on appropriateness; different condition values vary the levels of appropriateness. We observed a similar effect for aims and consequences. In this sense, beyond the specific methodological contributions this gauging introduces, the design also offers a path toward overarching conceptual questions regarding norm formation. Through meta-analysis of cases structured through this approach, it would be possible to better understand privacy norm formation across contexts.

9.4.2 *Practical*

The GKC-CI method is useful in emerging contexts, such as IoT, which often lack established norms. We first identify the various exogenous variables that act as a proxy to understanding respondents' disposition towards privacy. For example, certain respondents tend to be concerned about privacy and are actively pursuing ways to improve it for themselves. They read privacy policies, disable third party sharing, and find ways to circumvent the system whenever possible. Our CI analysis of the flows they deem acceptable confirms it: on average they tend to disallow flows, with notable exceptions when specific conditions, aims, and consequences align with social expectations. Another community perceives the polar opposite. They rarely read privacy policies, embrace third party sharing and

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

don't disable the tracking functionalities – all in the name of convenience and efficiency.

Furthermore, many respondents across interest groups perceive "anonymity" to be ineffective governance of information flows. "Anonymity" thus further fragments the overarching community of IoT users. In contrast to "consent," "anonymity" modifies information, rather than flow, impacting the association between information type and subject. Results indicate that adding "anonymity" as governance does not meaningfully impact perceptions of acceptability or behaviors.

Our results illustrate that governance of IoT should necessarily specify all parameters of the CI framework in structuring information flows, with clear identification of aims and conditions in the transmission principles. Practically, this means that when introducing new technology, it is possible to gauge the various factors using our methodology to reveal factors that have an effect on the acceptability of newly generated flows.

Furthermore, our results confirm previous findings that respondents (n=159) look for privacy policies to understand the privacy implications (e.g., Martin and Nissenbaum, 2016), however, some indicated in their comments that privacy policies are difficult to comprehend. Online forums and discussion with family were the other leading responses.

This result has practical implications with respect to how privacy related information could be structured and communicated so that users more intuitively understand. We propose that IoT manufacturers should clearly define all parameters according to CI and include institutional components within the transmission principle when prescribing information transfers. This could also offer a more informative and constructive discussion on the forums, with all the parameters stated explicitly.

9.5 CONCLUSION

We live in an age of great innovation! In the blink of an eye, information packets traverse the world; with a click of a button, information reaches millions of people. Things evolve at great speed and we, as a society, are looking for ways to keep apace with it. This forces us to adapt to the new reality and reconsider established concepts, such as the notion of privacy.

The GKC-CI method builds on the strength of two privacy theories. We use GKC to describe rules specific to a given context (rules-on-the-books and rules-in-use) and to understand users' strategies and norms. We use CI to gauge the appropriateness of information flows resulting from existing practices (rules-in-use) and/or prescribed by policy (rules-on-the-books).

Our results show diversity in respondents' privacy understanding and expectations around IoT devices. By gauging the information flows resulting from various practices employed by the Internet-connected systems, we can further see the

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

importance of contextual elements to gain deeper insights into their appropriateness. Specifically, we use the expressive language of GKC to further describe CI transmission principles. Results from survey questions that addressed CI and institutional aspects illustrate how more detailed conceptualizations of transmission principles, deconstructed using the attributes within the institutional grammar, highlight what aspects yield differences in respondents' opinions of information flows. This in turn helps to illuminate how particular aspects of institutional governance improve perceptions of these information flows to engender trust in governance.

ACKNOWLEDGMENTS

This project was supported by National Security Agency (NSA) Award (#H98230-18-D-006). Sincere thanks to Helen Nissenbaum and Jessie Garrett Taft for their help in finalizing, funding, and deploying the survey.

REFERENCES

Apthorpe, Noah, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. "Discovering Smart Home Internet of Things Privacy Norms Using Contextual Integrity." *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, no. 2 (2018): 59.

Bhatia, Jaspreet, and Travis D. Breaux. "Semantic Incompleteness in Privacy Policy Goals." In *2018 IEEE 26th International Requirements Engineering Conference (RE)*, pp. 159–169. IEEE, 2018.

Borkowski, Stephen, Carolyn Sandrick, Katie Wagila, Carolin Goller, Chen Ye, and Lin Zhao. "Magicbands in the Magic Kingdom: Customer-Centric Information Technology Implementation at Disney." *Journal of the International Academy for Case Studies* 22, no. 3 (2016): 143.

Crawford, Sue ES and Elinor Ostrom. "A grammar of institutions." *American Political Science Review* 89, no. 3 (1995): 582–600.

Frischmann, Brett M., Michael J. Madison, and Katherine Jo Strandburg, eds. *Governing knowledge commons*. Oxford University Press, 2014.

Geeng, Christine and Franziska Roesner. "Who's In Control?: Interactions In Multi-User Smart Homes." In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, p. 268. ACM, 2019.

Gorham, Ashley E., Helen Nissenbaum, Madelyn R. Sanfilippo, Katherine Strandburg, and Mark Verstraete. "Legitimacy in Context." At Privacy Law Scholars Conference (PLSC), University of California-Berkeley, 2019.

Joh, Elizabeth E. "The new surveillance discretion: Automated suspicion, big data, and policing." *Harv. L. & Pol'y Rev*. 10 (2016): 15.

Kumar, Deepak, Kelly Shen, Benton Case, Deepali Garg, Galina Alperovich, Dmitry Kuznetsov, Rajarshi Gupta, and Zakir Durumeric. "All Things Considered: An Analysis of IoT Devices on Home Networks." In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pp. 1169–1185. 2019.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press

Manikonda, Lydia, Aditya Deotale, and Subbarao Kambhampati. "What's up with Privacy?: User Preferences and Privacy Concerns in Intelligent Personal Assistants." In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 229–235. ACM, 2018.

Martin, Kirsten. "Information technology and privacy: conceptual muddles or privacy vacuums?." *Ethics and Information Technology* 14, no. 4 (2012): 267–284.

Martin, Kirsten and Helen Nissenbaum. "Measuring privacy: an empirical test using context to expose confounding variables." *Colum. Sci. & Tech. L. Rev*. 18 (2016): 176.

Nissenbaum, Helen. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.

Okoyomon, E., Samarin, N., Wijesekera, P., Elazari Bar On, A., Vallina-Rodriguez, N., Reyes, I., … & Egelman, S. (2019). On The Ridiculousness of Notice and Consent: Contradictions in App Privacy Policies. Privacy Law Scholars Conference (PLSC 2019), University of California, Berkeley, May 30–31, 2019. https://blues.cs.berkeley.edu/wp-content/uploads/2019/05/conpro19-policies.pdf

Sanfilippo, Madelyn, Brett Frischmann, and Katherine Strandburg. "Privacy as Commons: Case Evaluation Through the Governing Knowledge Commons Framework." *Journal of Information Policy* 8 (2018): 116–166.

Selbst, Andrew D. "Contextual expectations of privacy." *Cardozo L. Rev*. 35 (2013): 643.

Shvartzshnaider, Yan, Madelyn Sanfilippo, and Noah Apthorpe. "Privacy Expectations In the Wild: Integrating Contextual Integrity and Governing Knowledge Commons for Empirical Research." http://ssrn.com/abstract=3503096

Shvartzshnaider, Yan, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. "Learning privacy expectations by crowdsourcing contextual informational norms." In *Fourth AAAI Conference on Human Computation and Crowdsourcing*. 2016.

Tanczer, Leonie, Isabel Lopez Neira, Simon Parkin, Trupti Patel, and George Danezis. "Gender and IoT Research Report." University College London, white paper (2018): www.ucl.ac.uk/steapp/sites/steapp/files/giot-report.pdf

Valdez, André Calero and Martina Ziefle. "The users' perspective on the privacy-utility trade-offs in health recommender systems." *International Journal of Human-Computer Studies* 121 (2019): 108–121.

Zimmer, Michael. "Addressing Conceptual Gaps in Big Data Research Ethics: An Application of Contextual Integrity." *Social Media+ Society* 4, no. 2 (2018): https://doi.org/2056305118768300AU: DOI no?

Zou, Yixin, Abraham H. Mhaidli, Austin McCall, and Florian Schaub. " 'I've Got Nothing to Lose': Consumers' Risk Perceptions and Protective Actions after the Equifax Data Breach." In *Fourteenth Symposium on Usable Privacy and Security ({SOUPS} 2018)*, pp. 197–216. 2018.

https://doi.org/10.1017/9781108749978.010 Published online by Cambridge University Press