



A Roadmap for Applying the Contextual Integrity Framework in Qualitative Privacy Research

PRIYA C. KUMAR*, Pennsylvania State University, USA

MICHAEL ZIMMER, Marquette University, USA

JESSICA VITAK, University of Maryland, USA

Privacy is an important topic in HCI and social computing research, and the theory of contextual integrity (CI) is increasingly used to understand how sociotechnical systems—and the new kinds of information flows they introduce—can violate privacy. In empirical research, CI can serve as a conceptual framework for explaining the contextual nature of privacy as well as an analytical framework for evaluating privacy attitudes and behaviors. Analytical applications of CI in HCI primarily employ quantitative methods to identify appropriate information flows but rarely engage with the full CI framework to evaluate such flows. In this paper, we present a roadmap to guide HCI and social computing researchers on how to apply the full CI framework to qualitative projects. To help researchers envision what such an analysis can look like, each step includes an example analysis using interview data from projects on privacy and fitness tracking. We conclude by discussing how harnessing the full CI framework can address critiques of CI and identify opportunities for further theory development.

CCS Concepts: • **Security and privacy** → **Social aspects of security and privacy**; • **Human-centered computing** → Collaborative and social computing theory, concepts and paradigms

Additional Key Words and Phrases: Contextual integrity, privacy, information flows, theory, qualitative data analysis, methodology

ACM Reference format:

Priya C. Kumar, Michael Zimmer, and Jessica Vitak. 2024. A Roadmap for Applying the Contextual Integrity Framework in Qualitative Privacy Research. *Proc. ACM Hum.-Comput. Interact.*, 8, CSCW1, Article 219 (April 2024), 29 pages, <https://doi.org/10.1145/3653710>

1 INTRODUCTION

As social practices and digital technologies have become increasingly intertwined, public and academic interest in privacy has exploded [42]. Importantly, older theories of privacy—that predate social media and big data and focus on individual control over personal information [52]—no longer hold in an era when data collection is automated and ubiquitous and data is managed by individuals, other people, and corporations. Researchers in recent years have advocated for

*Corresponding Author

Authors' addresses: P.C. Kumar, priya.kumar@psu.edu, Pennsylvania State University, University Park, PA, USA; M. Zimmer, michael.zimmer@marquette.edu, Marquette University, Milwaukee, WI, USA; J. Vitak, jvitak@umd.edu, University of Maryland, College Park, MD, USA.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

© 2024 Copyright is held by the owner/author(s).
2573-0142/2024/4 – 219. <https://doi.org/10.1145/3653710>

frameworks and approaches that account for the networked and contextual nature of data privacy [58–60, 97].

The theory of contextual integrity (CI), which treats privacy as the appropriate flow of information within particular contexts [64, 66], is a popular approach for understanding and examining the privacy implications of sociotechnical systems [7, 11, 96, 97]. Compared to older privacy theories—which often frame privacy as something people either do or do not have—CI contends that privacy violations arise when information flows in a way that deviates from the norms of a given context. CI is a particularly useful framework for researchers to explore the nuances of data flows and the contextual factors that influence information disclosures and privacy concerns—something that is not addressed in detail by other theories. It identifies five parameters that shape how information flows and offers a framework for evaluating potential privacy violations. Thus, CI presents a philosophically grounded conception of privacy and a systematic process for analyzing privacy [64]. The strength of CI lies in its ability to translate the abstract notion of privacy into a concrete, observable unit of study [10, 11], and a growing community of scholars is working to extend its academic contributions. A decade ago, Barkhuus [10] called on the HCI and social computing communities to engage CI theory in their research on privacy. CSCW 2018 included a workshop focused on incorporating the CI framework into empirical research [8] and the annual Symposium on Applications of Contextual Integrity is entering its sixth year.¹

In empirical research, CI can serve as a conceptual framework for explaining what privacy is, as well as an analytical framework for investigating how privacy manifests in technologically mediated interactions. HCI and social computing researchers often apply CI as a conceptual framework, meaning scholars reference CI’s focus on “context” when defining privacy or motivating their research questions. Less often, researchers use CI to inform the design of their study or the analysis of their data. (See [7] for an early review of CI use in HCI studies.) When researchers do take up CI as an analytical framework, they primarily apply it via quantitative methods (i.e., surveys) and engage only with the descriptive portion of the framework, which examines how information flows align with context-relative norms [7, 11]. While these applications of CI in prior work provide useful contributions to privacy research, we argue that the narrow focus of this work presents an opportunity for researchers to also harness the prescriptive portions of the CI framework and integrate broader socio-political considerations into their privacy analyses—something qualitative researchers are well-equipped to achieve.

Therefore, this paper provides qualitative researchers with a roadmap for integrating the full CI framework into their data analysis. To do this, we first describe the core components of CI, then provide a detailed description of Nissenbaum’s nine-step decision heuristic for determining whether a new information practice violates privacy norms, and if so, how to respond. The first five steps of the heuristic are *descriptive* in that they ask the researcher to identify and define the information practice of study, the parameters associated with that practice, and the norms guiding appropriate information flows. The final four steps shift from descriptive to *prescriptive* and ask researchers to evaluate their findings from the initial steps in light of wider social, moral, and political values and determine whether the practice should continue. To complement this roadmap and illustrate how it can be applied in qualitative research, we provide an example analysis for each step, focusing on the flow of personal fitness information (PFI) from wearable fitness trackers into the health care context. While the roadmap is intended to be customized to

¹ See <https://privaci.info>

a given topic or question of study, we hope this example analysis helps researchers better understand how to engage the full CI framework in their own work.

This paper offers important theoretical and methodological contributions for qualitative privacy researchers who have selected CI as the theoretical framework to guide their analysis.² We provide researchers with a step-by-step process to apply CI as an analytical framework in their studies, including questions to guide them through each step in the process. We highlight how deeper engagement with the framework enables researchers to identify when practices should be stopped or re-evaluated and when privacy-violating practices can be acceptable (e.g., when they are balanced against a greater social good). Finally, we offer detailed guidance explaining two frequently misunderstood aspects of CI—context and transmission principles [66]—and discuss several avenues through which to further extend the theory.

2 BACKGROUND

In this section, we summarize the main tenets of CI and review how existing HCI and social computing research has used CI to study questions about privacy and technology. We also explain what the community stands to gain by employing CI as a method for analyzing qualitative data.

2.1 Overview of Contextual Integrity (CI)

Privacy has increasingly been understood as a process of managing boundaries across different spheres—boundaries that shift as contexts change and as new technologies disrupt them [71]. Such dynamism is central to the notion of networked privacy, which Marwick & boyd define as the “ongoing negotiation of contexts in a networked ecosystem in which contexts regularly blur and collapse” [58:1063]. Wu and colleagues build on this work and encourage researchers to embrace “an even broader contextual view of privacy” [97:3] that spans individuals, groups, and society.

Contextual integrity is a conceptual framework that takes context as its starting point. The framework rests on the understanding that social interactions occur in particular contexts, and that norms govern our expectations of how information should flow in a given context. As such, CI rejects the public/private dichotomy, tying adequate privacy protection to the value of respecting informational norms within specific contexts [64]. Thus, CI provides a helpful framework for explaining why certain information flows are acceptable in one context but problematic in another.

CI was not developed to be a grand theory of privacy and why it matters. Rather, CI responds to the observation that “novel sociotechnical practices” often provoke privacy concerns [66:224]. More specifically, Nissenbaum developed CI “in an attempt to understand what people saw threatened by novel sociotechnical practices wrought by a family of technologies, including computers, digital networks, information systems, databases, communications media, electronic hardware, and software” [66:224]. This focus on the sociotechnical dimension of privacy is what makes CI especially well-suited to privacy research in the HCI and social computing space. The building block of CI is the concept of an “information flow,” which refers to “the passage or transmission of information or data from party (or parties) to party (or parties)” [66:225]. The premise underpinning CI is that actions, behaviors, and practices generate information flows. CI further recognizes that information flows about people are what drive a “robust social sphere” [66:225]. In other words, society needs information about people to circulate, and society is richer

² For information about other privacy theories that qualitative researchers can use to guide their work, see Wisniewski and Page [96] and Margulis [52].

when that happens. CI posits that privacy is not the absence of information flow, but rather the presence of appropriate information flow.

To determine what constitutes an appropriate flow, CI turns to the concept of norms. It argues that flows of information are appropriate when they conform to informational norms “that describe, prescribe, proscribe, and establish expectations for characteristic contextual behaviors and practices” [66:227]. The task of a CI analysis is to define the sociotechnical practice in question, identify whether its information flows map onto context-specific norms, and offer recommendations about what, if anything, should be done to address any concerns about the information flow.

2.1.1 The Contextual Integrity Decision Heuristic. To help people apply CI to specific practices, Nissenbaum [64] offers a nine-step decision heuristic. Since privacy, in CI terms, is the extent to which information flows in a manner that aligns with the expectations of a given context (i.e., context-relative informational norms), Steps 1-4 of the heuristic involve identifying the parameters that comprise context-relative informational norms. Steps 5-8 evaluate the conflicts, consequences, and implications of technologically driven changes to these norms. Step 9 involves making recommendations about how to respond to a new technology or practice. In Section 3, we explain each step in detail and offer guidance on how researchers can apply the heuristic to their studies.

The CI heuristic can be used “as a descriptive tool, systematically accounting for people’s reactions to the myriad technical systems radically affecting flows of personal information” and as a prescriptive tool “for evaluating these systems from a moral and political point of view” [64:189]. But, as others have noted, scholarly applications of CI in HCI and social computing have primarily engaged with CI’s descriptive capacities [7, 11]. We echo Badillo-Urquiola et al. [7], Benthall et al. [11] and Mir [61] in encouraging scholars to harness the descriptive *and* prescriptive capacities of CI to address the privacy issues facing society. We also recognize that the full CI framework can be intimidating, so this paper provides a roadmap to help researchers better understand how they can apply CI to their own work. But first, we review how CI has been taken up in HCI and social computing, differentiating between conceptual and analytical applications of the framework.

2.2 Research Applications of Contextual Integrity

Scholars across a range of disciplines have employed CI to analyze legal doctrine [40] and information governance [79], interpret legal frameworks [32], inform organizational decision-making [35], design computational systems [11], measure people’s privacy preferences [54], and explore privacy in practice [15]. HCI and social computing scholars use CI as a conceptual framework to understand privacy issues, and, to a lesser extent, as an analytical framework to examine empirical data. We highlight these uses below.

2.2.1 Conceptual Applications of Contextual Integrity. Conceptually, scholars have engaged CI to consider privacy issues related to search engines [101], social media [36, 47], location sharing [10], facial recognition systems [67], COVID-19 contact tracing applications [93], and sociotechnical systems more broadly [61]. These non-empirical articles use CI to explain how engagement with sociotechnical systems raises privacy concerns and how such concerns should be addressed.

CI has also been used as a conceptual framework in empirical studies, with Barkhuus [10] encouraging the HCI and social computing community to use it in empirical analyses of privacy because it provides a vocabulary for nuanced insights into how and why people respond the way they do when engaging with sociotechnical systems. Since then, scholars employing a variety of methods have incorporated CI into their studies, largely by linking their findings to specific aspects of CI. For example, one team studying people's responses to unexpected smartphone data flows used CI to explain perceptions of "creepiness" [81].

Researchers conducting experiments and testing novel systems have used CI to develop hypotheses about how people respond to targeted advertising [33] and automated inference of personality traits [31]. Others have used CI's conception of norms to understand participant perceptions of Internet of Things devices [2] and contact tracing applications [37]. More specifically, researchers have identified how the CI parameters of actors, information types, and transmission principles shape norms surrounding smart homes [99], ephemeral social media [98], and digital COVID-19 certificates [68]. Finally, research teams have also used empirical analyses to develop new privacy-related conceptual frameworks that build on and extend CI [18, 27]. These studies, while empirical, still engage CI on a conceptual, rather than analytical, level. That is, these studies link their findings to CI, but they do not use CI to design their study or analyze their data directly.

2.2.2 Analytical Applications of Contextual Integrity. Researchers are beginning to harness the analytical power of CI, but primarily via quantitative methods. Technology ethics scholar Kirsten Martin [53] adapted the factorial vignette survey methodology to measure privacy norms as defined by CI. In this methodology, researchers present participants with a series of brief vignettes and ask for their responses. Each vignette contains a set of factors, in this case, CI parameters, which are randomly changed as the vignettes cycle through. For instance, a sample vignette from Vitak et al.'s [91] study comparing privacy attitudes in the US and Netherlands was: "Instagram acquires one year's worth of your physical activity (inferred from phone stats). They plan to use this data to infer your political views with the goal of creating a national database of citizens," with the factors underlined. Participants responded to two Likert-scale statements per vignette: "This use of my data is appropriate" and "This use of my data would concern me." Researchers can statistically analyze differences in responses to understand how changes in parameters affect privacy perceptions and norms.

Martin and Nissenbaum refined their application of CI to the factorial vignette methodology to study privacy norms related to general information flows [54], public records [55], and location data [56]. This operationalization of CI has taken root in HCI and social computing, with scholars using the methodology to study privacy expectations and norms related to general data flows [91], smart home technologies [1, 4, 82], smart toys and their alignment with privacy law [5], health data [28], social media data [30] and COVID-19-related systems [90, 100]. Researchers have also used the CI-based factorial vignette approach to develop a methodology for automating the process of identifying norms [83], demonstrating CI's promise as an analytical tool for shaping research design itself.

The few qualitative applications of CI as an analytical framework have largely involved coding textual data (e.g., from interviews, focus groups, online comments, observational field notes, etc.) and interpreting it using CI, organizing the findings around CI's concepts, and linking CI to the study's broader contribution [12, 15, 45, 46, 80]. What differentiates this work from conceptual applications of CI in qualitative research is the degree to which the study integrates CI. In analytical applications, researchers use CI to analyze their data and embed CI throughout their

study’s findings and contributions. For example, in their work on children’s understandings of privacy online, Kumar et al. [45, 46] qualitatively coded interviews with families and then analyzed the coded data for information related to CI’s parameters. They organize their findings around contexts and CI’s parameters and then use their findings to explain how educational efforts can incorporate elements of CI to strengthen children’s privacy literacy. Similarly, in a study on privacy in citizen science initiatives, Bowser et al. [12] qualitatively coded their interview and focus group data using codes related to CI and organized their findings around CI concepts. Their analysis is notable for its engagement with the second half of the CI framework, which addresses how social values inform judgments of appropriate information flow.

In sum, methodological innovation in CI’s analytical capabilities has so far concentrated on quantitative research approaches and the descriptive dimension of CI. For privacy researchers wanting to explore the nuances of qualitative data, we offer the following roadmap to help them embed CI into their research in a way that harnesses the full analytical power of the framework.

3 A ROADMAP FOR ANALYZING QUALITATIVE DATA THROUGH THE CONTEXTUAL INTEGRITY FRAMEWORK

In this section, we review each step of Nissenbaum’s CI decision heuristic and offer guidance on how researchers can apply each step to analyze their qualitative data. Table 1 synthesizes this guidance by identifying what each step is, what it aims to accomplish, and what guiding questions researchers can consider to determine how to apply that step to their analysis. To help researchers envision what this process looks like, we also provide an example in each step that highlights one approach to applying the CI framework with qualitative data.

Table 1. Guidance for Applying the Nine-Step CI Decision Heuristic to Qualitative Data Analysis

#	Step (Text from [64:182])	Goal/Aim	Guiding Questions
1	Describe the new practice in terms of information flows.	To define the object of study in the language of CI.	• In this study, what are people, groups, or institutions doing? • How are they doing it (e.g., what tools are they using?) • How does information fit into this practice? • What aspects of the practice are novel, and what are extensions of existing practices?
2	Identify the information types, activities, and purposes involved in a given information flow and link them to a prevailing context.	To embed the information flow in a social context.	• What types of information does the flow involve? • What actions are occurring in this practice? • Why is this information flow happening? • What are actors using the information for?

3	Identify information subjects, senders, and recipients.	To clarify who and/or what is participating in the information flow.	• Who/what is sending information? • Who/what is receiving information? • Who is the information about, or to whom does the information pertain?
4	Identify transmission principles.	To establish the conditions that govern the information flow.	• What requirements need to be met for this information flow to occur? • In what circumstances can the information flow occur or not? • What stipulations exist that dictate whether this information flow can occur?
5	Locate applicable entrenched informational norms and identify significant points of departure.	To discern expectations about how information usually flows in the context being studied and to pinpoint what the new practice changes about information flows.	• How is information typically used and managed in this context? • How do these uses of information align with the broader goals, values, or purposes of the context? • What, if any, specific aspects of the information flow are altered by the practice under study?
6	Prima facie assessment of contextual integrity.	To determine whether the practice under study violates privacy.	• Does the information flow align with entrenched norms of this context, as established in Step 5? Why or why not?
7	Consider moral and political factors affected by the practice in question.	To recognize the social implications of the practice under study.	• How does the practice under study threaten autonomy or freedom? • How does the practice alter power structures or power relations? • How does the practice affect equality, fairness, justice, democracy?
8	Consider the meaning or significance of moral and political factors in light of contextual values, ends, purposes, and goals.	To assess how the practice under study could affect the context.	• How does the practice under study align with the aims or goals of the context? • How do the moral/political implications of the practice under study advance or undermine these contextual aims?
9	Recommendation for or against the system or practice being studied.	To judge, based on the analysis, what the course of action should be.	• Should the practice under study continue as is, or should it be rejected? • What, if any, modifications or conditions should be implemented?

The example analysis draws on research data from interviews with 54 people across two projects on privacy and fitness tracking. The first project included 21 people interviewed in person from March to July 2013 [73, 74]. The original researchers shared the anonymized interview transcripts and field notes with the authors for analysis in accordance with the project's ethical approvals. The second project included 33 people interviewed by the authors from March to April 2017 [103]. The interview procedures from 2013 and 2017 included similar questions about how interviewees used Fitbit and how comfortable they felt sharing personal fitness information (PFI) with various actors. Our example draws on the subset of data from each project on PFI flows pertaining to health care, since wearable fitness trackers are often marketed as ways for people to improve their health [17]. This example analysis is not intended to provide an empirical contribution nor a temporal comparison, but rather to walk through the roadmap using actual data on a relevant privacy issue, so that researchers can gain an understanding of what kind of information to look for in their own data to address the guiding questions. Our aim is to demonstrate how researchers can use the CI framework to analyze their qualitative data, and we believe this roadmap can be applied to a variety of qualitative datasets and privacy-related topics.

Step 1: Describe the new practice in terms of information flows.

The first step of the decision heuristic asks researchers to define their object of study in the language of CI. Researchers can do this by considering the following questions, connecting responses to prior work as applicable: *In this study, what are people, groups, or institutions doing? How are they doing it (e.g., what tools are they using)? How does information fit into this practice? What aspects of the practice are novel, and what are extensions of existing practices?* Researchers may not be able to answer all of these questions before data analysis; indeed, these questions may be the focus of the study. But considering these questions at the start of the analysis can help researchers better grasp what they are examining and what CI offers to the analysis, even if that object of focus shifts over the course of the study.

Step 1 Example: The interview data focuses on the practice of self-tracking personal fitness information (PFI) using Fitbit devices [17]. Wearable fitness tracking is increasingly popular. In 2012, 8 percent of U.S. adults used a mobile app or online tool to track health or fitness-related information [26]; by 2020, 21 percent of U.S. adults regularly wore a wearable fitness tracker or smartwatch [94]. When an individual wears a fitness tracker, the device generates various kinds of data about the person's body and physical activity (e.g., heart rate, steps taken). The specific makeup of an individual's PFI varies based on the type of device they use, what types of information they track, and to what extent they link their device to other fitness-related apps or services. The device syncs with a cloud-based platform (e.g., Fitbit Dashboard), from which the individual can view the data. The platform may also include a social element (e.g., Fitbit Community) where individuals can connect with other users and share information about their performance or achievements. While the core practice—self-tracking quantified information about the body—is not new [95], the sociotechnical system in which the practice is embedded (e.g., a wearable device connected to a digital platform and online community) yields novel information flows. **A CI analysis can identify what, if anything, about the practice of Fitbit-based self-tracking may pose privacy concerns, and if so, how society should respond. To manage the scope**

of the example analysis, we focus this example on the privacy implications of PFI flows in health care.

Step 2: Identify the information types, activities, and purposes involved in a given information flow and link them to a prevailing context.

This step asks researchers to embed the information flow in a social context. As the name “contextual integrity” suggests, context comprises the heart of CI. But the term “context” is multifaceted. When it comes to information flows, context can refer to a particular platform or system (e.g., Facebook), industry or sector (e.g., financial services), business model or practice (e.g., marketing), or social domain (e.g., education). The CI framework approaches context as a social domain, contending that when technologies introduce new information flows, an understanding of context as social domain is most useful when determining whether such flows serve the best interests of society [65].

It is important to avoid conflating context with place [66]. Although certain places, like hospitals and schools, are often tightly linked to specific contexts, like health and education, respectively, these links are not a given. For instance, a teaching hospital is also part of an educational context, and teachers experience a school within an employment context. Furthermore, a place-based understanding of context limits context to serving as a container for social action. However, “[r]espective roles, activities, purposes, information types do not exist in a context; rather, these factors *constitute* a context” [66:227, emphasis in original]. Contexts do not necessarily exist a priori; rather, they emerge from a confluence of existing factors, including the broader social, cultural, and political environments in which technologies exist. Hence, the researcher’s task is to identify the elements that a context brings together.

Researchers can do this by considering the following questions: What types of information does the flow involve? What actions are occurring in this practice? Why is this information flow happening? What are actors using the information for? From these questions, researchers can infer the activities and purposes embedded in the practice. An activity is descriptive: for instance, a health insurance company might receive step data from a workplace wellness program. A purpose is explanatory: an insurance company might use that step data to offer premium discounts to customers. Nissenbaum [66] has acknowledged that early work on CI insufficiently explained how factors like “purpose” fit into the framework. Recent work has situated purpose within transmission principles, the focus of Step 4 [27, 82]. However, we situate purpose with context because, as Nissenbaum states, “[s]ocial contexts are what they are because of respective contextual aims, purposes, and values.” [66:227]. We believe that researchers can identify what purposes align with a given context in Step 2, and that they can consider how information flows align with those purposes in Step 4.

To identify context, researchers can review the answers to these questions and identify meaningful overlaps between information types, activities, and purposes. Depending on their study design and research questions, researchers may find one overarching context, multiple distinct contexts, or multiple overlapping contexts. It is important to keep in mind that contexts do not operate universally. For instance, most people in the U.S. obtain health insurance through their employer [39], which creates overlaps in the healthcare and employment contexts that may not exist in other countries. Focusing on contexts as configurations of several elements can attune researchers to the specific ways that context manifests in their study.

Step 2 Example: Wearable fitness trackers generate PFI, whose information types include a user's step count, heart rate, calories burned, and the frequency and duration of physical activities and sleep [17, 75]. Since this example analysis focuses on the privacy implications of PFI flows into health care, we consider the elements that constitute the health care context. Interviewees noted that they often share information resembling PFI with doctors during appointments, and that doctors may obtain such information through tests. One described that "doctors all have laptops now when you walk in and they're inputting info. They seem to collect more data today than they did years ago," an activity likely linked to the rise of electronic medical records in health care [15]. Interviewees expressed that doctors use this information for the purposes of monitoring patients' health and offering guidance or advice to treat medical problems. They also recognized that their information may flow beyond doctors and to insurance companies, who may use it for the purposes of processing claims and setting coverage or policy rates. Even this brief description of purposes points to different ends: providing care and paying for it. Nissenbaum [64, 66] notes that the elements that constitute a context—including functions, purposes, and values—are contentious and far from clear cut. In health care, the value of care may exist in tension with the value of efficiency. This will be important for the normative portion of the CI analysis, especially Step 8.

Step 3: Identify information subjects, senders, and recipients.

This step asks researchers to clarify what actors are participating in the information flow. Actors can be "single individuals, multiple individuals, or even collectives such as organizations, committees, and so forth" [64:141]. CI includes three kinds of actors. Senders are the entities that transmit information. Recipients are the entities to whom information is transmitted. Subjects are those to whom the information pertains. Researchers can identify them by considering the following questions: *Who/what is sending information? Who/what is receiving information? Who is the information about, or to whom does the information pertain?* Given that information flows can involve myriad actors, researchers may need to narrow their focus to only a few, depending on their research questions.

Step 3 Example: With wearable fitness trackers, the subject and the sender are often the same: a Fitbit device generates data pertaining to a given user, and that user makes decisions about sending the information to a particular recipient. However, two interviewees noted that someone could put their Fitbit on a dog or give the device to their child to play with. Thus, fitness tracker data could have multiple *subjects*. This has implications for the accuracy of the data, something we discuss in Step 6. Others noted that the company Fitbit could disclose a user's PFI, illustrating a case where the *sender* (Fitbit company) and the *subject* (Fitbit user) are different. Indeed, one reason that wearable fitness tracking has raised privacy concerns is the fact that data often flows to many recipients [17], including medical researchers, insurance companies, and employers [14, 23, 63, 78]. Since this example focuses on the privacy implications of PFI flows in health care, we concentrate the analysis on interviewees' responses to questions about PFI flows to two particular recipients: doctors and insurance companies.

Step 4: Identify transmission principles.

This step asks researchers to establish the conditions that govern the information flow. In CI, these conditions are called transmission principles. The transmission principle is one of the most noteworthy elements of the CI framework, but also one of its most ambiguous. It refers to a “constraint on the flow (distribution, dissemination, transmission) of information from party to party in a context...express[ing] terms and conditions under which such transfers ought (or ought not) to occur” [64:145]. For instance, when one friend says to another, “Don’t tell anyone else—I’m pregnant,” the speaker is requesting the recipient keep the information to themselves and prevent it from flowing to another recipient. This constitutes a transmission principle of secrecy.

Transmission principles can be explicit (i.e., codified in law, policy, or direct statements, like the one in the previous paragraph) or implicit (i.e., implied or inferred based on the situation at hand). Examples of explicitly established transmission principles include a law requiring doctor-patient conversations to be confidential, a privacy policy stating that a company will only disclose information with an individual’s consent, or a convener stating that a meeting is operating under the Chatham House Rule³. An example of an implicit transmission principle would be a person who shares information about a miscarriage with a close friend and assumes the details will stay confidential between the pair.

Nissenbaum describes the transmission principle as perhaps “the most distinguishing element of the framework of contextual integrity; although what it denotes is plain to see, it usually goes unnoticed” [64:145]. However, “many have found the transmission principle (TP) parameter to be puzzling because, on its face, it is less familiar to accounts of privacy than actor-capacities and information types” [66:230]. There is no defined list of transmission principles; indeed, Nissenbaum notes that “[t]he list is probably infinite” [64:145] given the variety of circumstances in which information flows occur. Prior work has identified several transmission principles, including confidentiality, desert (deserving to know), entitlement, compulsion, need, voluntary, notice, consent, exchange [64], reciprocity, anonymity [12], temporality [15], mutuality, requirement, and secrecy [45]. Researchers may wish to take these as a starting point while also heeding Nissenbaum’s [66] argument that an endless array of transmission principles can govern information flows.

Researchers can identify transmission principles by considering the following questions: What requirements need to be met for this information flow to occur? In what circumstances can the information flow occur or not? What stipulations exist that dictate whether this information flow can occur? The process of identifying transmission principles may be more deductive or inductive, depending on a study’s research questions and goals. If transmission principles are codified, such as a legal requirement for consent, researchers may examine whether and how this requirement is taken up in practice. Conversely, an exploratory project may involve inductively analyzing data to distill what transmission principles govern flows in practice.

Step 4 Example: Since transmission principle can be hard to grasp, we offer the following method as one way to identify transmission principles in text-based data (e.g., interviews, documents, field notes). First, we located all quotes in the interview data that expressed the sentiment, “This information flow is fine IF...” The conjunction “if” signaled that support for the information flow was contingent on it meeting a given condition.

³ See <https://www.chathamhouse.org/about-us/chatham-house-rule>

Identifying the condition would yield the transmission principle at play. We looked for patterns, grouping quotes that expressed a similar condition. We then examined each group to identify a common characteristic, from which we discerned the transmission principle.

Table 2 lists the transmission principles we identified in the data, along with sample quotes. This list is not a definitive list of transmission principles in CI; recall Nissenbaum’s statement that transmission principles are infinitely variable [64]. Rather, this list illustrates one way researchers can qualitatively analyze text-based data to identify transmission principles, something prior work has established as a challenge [38, 102]. This method can be compatible with approaches like Shvartzshnaider et al.’s, [82], which divides the transmission principle into the sub-parameters of aims, conditions, modalities, and consequences. Researchers could use these sub-parameters as categories as they group their units of data. For instance, the transmission principle of need listed in Table 2 may align with their sub-parameter of aim, since both reflect that an information flow may be permissible IF it is necessary to fulfill one of the goals of the context.

Table 2. Range of Transmission Principles Governing PFI Flows

Sentiment: “The information flow is fine IF…”	Transmission Principle	Sample Quote from Data
It is necessary to fulfill the goals of the context.	Need	“If my doctor specifically needs it for my care, then I would be willing to provide it of course.”
The subject chooses to disclose or transmit the information.	Voluntary	“If I chose to share it with my doctor, that’s okay. But I want that to be my decision.”
The subject is told how the information will be used and/or managed.	Notice	“I would like to have the doctor say, ‘Hey, this is why I want to do this, this is the benefit,’ and is the doctor giving that information to anybody? Where would it be going?”
The recipient asks the subject and/or the subject gives permission.	Consent	“If a doctor asks me, I will provide it. But I would like to personally provide it to them, not them asking Fitbit without my permission and just getting my information, even if it’s for my benefit”
The subject gets something in return.	Exchange	“I think the fitness data would have to be tied to something to share it, again, with some sort of incentive. Either they can help me out or decrease premiums or something like that. There’d have to be a clear benefit to me because otherwise you already know about all of my health. You are the insurer.”
The subject must disclose or transmit the information.	Mandatory	“Insurance companies, nope, unless somehow I’m required to tell them [how often I exercise].”

The information is not recorded.	Ephemerality	"I would be more comfortable with my doctor knowing it. I would [be] more uncomfortable with it being in a file."
The recipient wants the information.	Desire	"Doctor, I guess if they wanted to know for whatever reason they can."
The recipient does not disclose or transmit the information.	Secrecy	"I guess if it were my doctor and only my doctor."
The information does not identify the subject.	Anonymity	"Yeah, I'm more anonymous there. Like I hope they [insurance] wouldn't track me directly."
The information is not granular.	Aggregation	"I think it would be very useful to have the doctor able to see the activity, maybe not the activity but the summary."

Step 5: Locate applicable entrenched informational norms and identify significant points of departure.

This step has two components, asking researchers to (1) discern expectations about how information usually flows in the context being studied (i.e., norms) and (2) pinpoint what the new practice changes about information flows. Norms, which serve as the foundation of the CI framework, are the explicit or implicit rules that "describe, prescribe, proscribe, and establish expectations" for behavior in a given context [66:227]. Norms "may emanate from a variety of sources, may or may not be enshrined in law, may be commanded or merely emergent, may vary over time and across cultures, may be strict or approximate, may be universally or merely locally known, and so forth" [66:227]. Despite such heterogeneity, CI approaches norms as entrenched; in other words, norms "reflect a settled accommodation" [66:234]. Indeed, when people's preferences regarding specific information flows are surveyed, rather than their attitudes toward general concepts like privacy or control, similarities in judgments emerge [64:151].

CI recognizes that norms "govern the flow of personal information in distinct social contexts" and contends that "[i]nformation technologies alarm us when they flout these informational norms" [64:3]. In other words, privacy concerns arise out of violations of context-specific norms. But before researchers can determine whether an information flow poses privacy concerns—the focus of Step 6—they must establish what norms govern the practice under study. Researchers can do this by considering the following questions: *How is information typically used and managed in this context? How do these uses of information align with the broader goals, values, or purposes of the context? What, if any, specific aspects of the information flow are altered by the practice under study?*

Since CI is based on information flows, researchers should focus on the norms that govern the "transmission, communication, transfer, distribution, and dissemination" of information in a given context [64:140]. Researchers can look for information about norms in their data, but it is important to remember that the entrenched quality of norms implies that they exist within shared understandings of a particular context. Thus, it can be useful to link information about norms from one's data with information from other relevant domains, including law, culture, scholarship, and social practice.

Step 5 Example: Interviewees indicated that the act of patients discussing information about their exercise and sleep with doctors is not unusual, provided that such discussion is necessary for patient care and something patients understand and agree to. Framed in CI terms, this suggests that in the healthcare context, the information flow of patients (*senders*) choosing (*transmission principle*) to provide PFI (*information type*) about themselves (*subjects*) to doctors (*recipients*) who need it (*transmission principle*) is a fairly stable norm. Indeed, health communication literature supports this, stating that “[t]he communication of relevant information is a central element of health care” that enables people to monitor their own conditions and doctors to diagnose, monitor, and treat conditions in patients [43:238]. In CI terms, the flow of relevant information is essential to meet one of the purposes of health care—managing illness.

However, integrating the practice of wearable fitness tracking into this information flow presents several points of departure. First, wearable fitness tracking changes the nature of the information being discussed. Wearable fitness tracking tends to be automated and continuous, producing very granular data about one’s body and activities. Second, the entry of Fitbit as an intermediary in the practice of self-tracking introduces novel forms of information flow: people can disclose PFI (e.g., tell their doctor their step count), show PFI (e.g., pull up their Fitbit dashboard during a doctor’s appointment) or transmit PFI (e.g., sync a Fitbit with an electronic medical record system). Third, elements of the information transmission process can be automated. Fitbit users can configure their devices to automatically sync with Fitbit’s servers and link their Fitbit data with third parties [20].

The company Fitbit has sought to make Fitbit data easier to integrate with electronic medical records systems, insurance programs, and more [23, 24]. In CI terms, the information flows generated from the practice of wearable fitness tracking may include more detailed *information types*, more *transmission principles* for consideration (e.g., ephemerality, aggregation), and additional *recipients* of information (e.g., insurance companies, employers).

Step 6: Prima facie assessment of contextual integrity.

Step 6 marks a shift in analysis from descriptive to prescriptive. In Steps 1-5, researchers examine the information practice under study and identify the parameters and components pertinent to CI (i.e., context, actors, information types, transmission principles, and norms). Steps 6-9 ask researchers to connect these findings to broader social commitments and make a series of judgments, culminating in a recommendation about whether the practice under study should continue.

In this step, researchers take the findings from Step 5 and determine whether the practice under study violates privacy. As noted above, norms serve as the benchmark for judging whether the practice under study is appropriate, and privacy constitutes the appropriate flow of information. If an information flow aligns with the norms of a given context, it likely does not violate privacy because people experience the flow as something that makes sense. However, if one or more aspects of an information flow deviate from norms, the practice under study may violate privacy, since it does not conform to established expectations. In CI terms, “a practice violates a privacy norm if resulting flows fail to map onto expected values for the parameters” [66:231].

Researchers can draw this conclusion by considering the following questions: *Does the information flow align with entrenched norms of this context, as established in Step 5? Why or why not?* Researchers may find it easier to start by identifying what, if any, concerns regarding the practice under study exist in their data and then linking those concerns to specific CI parameters. If no concerns exist, or if those concerns do not pertain to CI parameters, the practice may align with entrenched norms and thus not violate privacy (though it could certainly pose other problems). If this is the case, researchers may not need to complete the remaining steps. Otherwise, the practice likely violates contextual integrity.

It is important to acknowledge that in CI, the mere existence of a privacy violation is not inherently problematic. In other words, the fact that a practice violates privacy is not sufficient grounds to dismiss the practice. The purpose of the next two steps is to determine whether the privacy violation is significant enough to warrant rejecting the practice.

Step 6 Example: The example analysis for Steps 3-5 suggests that the flow of PFI from wearable devices into health care can introduce points of departure in all five CI parameters: *subjects* and *senders* may be distinct, information may flow to additional *recipients*, the *type of information* generated is much more granular, and more *transmission principles* may need to be considered. Health care relies on timely and accurate information flows between patients and providers [43]. Thus, at first glance, the flow of PFI from wearable devices into this context would seem to be a boon. However, research has raised questions about the accuracy of data generated by wearable fitness tracking [19, 75], as did interviewees. For instance, one said they wore their device during a hurricane, and the rapid changes in air pressure led the Fitbit to mistakenly record them as having climbed dozens of flights of stairs. As a result, they questioned the utility of such data flowing into a healthcare context, comparing the unreliable nature of Fitbit data with a “fact” like an x-ray. Another expressed similar concerns regarding the accuracy of PFI flows and concerns with data being shared with an insurance company.

Beyond issues with the device itself, data accuracy could be compromised if someone other than the Fitbit user wears the device, as noted in Step 3. Such questions about accuracy suggest that the *automated* transmission of PFI to a doctor or insurance company could violate contextual integrity. Indeed, one interviewee responded to the idea of Fitbit data going straight into a medical file by saying, “The automatic aspect of that...seems odd.” Another added, “I feel like it wouldn't hurt to just have a normal appointment and talk about how active I am rather than see my Fitbit data.” One interviewee described their concern about PFI flows to insurance companies by comparing them to flows with doctors; because of the nature of the relationships, this person said they’d “feel more comfortable with it being a personal relationship....At least with the doctor you get to see and meet your doctor. You don't get to see and meet your insurance company so the doctor may understand the circumstances.”

As explained in Step 5, it can be considered normal for patients to willingly discuss PFI with doctors when such information is necessary. Sharing PFI with doctors through conversation gives patients a chance to put the information in context. But the introduction of device-driven information flows to doctors and health insurance companies could disrupt that balance. As noted in this step, people expect information used to make medical decisions to be as accurate as possible. Consumer devices like Fitbit do not necessarily meet

that standard, and automated transmission would deprive patients of an opportunity to participate in decision-making involving their own health. Thus, the automated transmission of wearable device-generated PFI for the purpose of providing health care could pose a *prima facie* violation of contextual integrity.

Step 7 (Evaluation 1): Consider the moral and political factors affected by the practice in question.

This step asks researchers to recognize the social implications of the practice under study. Where the previous step focused on identifying whether a privacy violation exists, this step focuses on what is at stake. Step 7 is where privacy's significance as a social value comes to the fore. Even if a study's data only focuses on individual consequences, the step's reference to moral and political factors asks researchers to consider effects on collective social functioning. Indeed, that is why the step itself is labeled as an evaluation.

Researchers can evaluate the social implications of the practice under study by considering the following questions: How does the practice under study threaten autonomy or freedom? How does the practice alter power structures or power relations? How does the practice affect equality, fairness, justice, democracy etc.? As with Step 5, researchers may find it useful to look for information in their data about the effects of the practice under study and then interpret that information using other sources or domains. These can include applicable laws, proposed policies, and/or relevant moral frameworks or political theories, depending on the research questions guiding the study. Researchers may find that turning to sources outside their data yields insights for this step. For instance, media coverage may document specific occurrences of benefit or harm arising from the practice under study; corporate documents may shed light on potential future directions for the practice under study; and scholarly literature may establish theoretical or conceptual connections between the practice under study and broader moral and political factors.

Step 7 Example: As noted in the Step 6 example, when a subject's PFI is automatically transmitted to a doctor or health insurance company, they have less opportunity to contextualize the data. PFI is often approached with "mechanical objectivity"—seen to accurately depict a subject's physiological state [72]. Yet subjects also develop a "situated objectivity" toward PFI, interpreting it as part of their lived experience, expectations, and cultural understandings [72]. Automated transmission of PFI to doctors can foreclose subjects' opportunities to bring their situated objectivity into decision-making concerning their health, which could reduce their autonomy in this context. Automated transmission of PFI to insurance companies, particularly when tied to incentives (e.g., discounts) or penalties (e.g., higher premiums), can be coercive, which Loi et al. [48] present as a moral wrong because it infringes on a subject's autonomy. Such infringement is particularly problematic considering the fact that PFI flows exacerbate information asymmetries between subjects and insurance companies, tilting the balance of power more heavily in favor of insurance companies [86]. This is because insurance companies have the expertise and incentive to identify patterns in a subject's data, compare those patterns against others (whose data they also manage), and use those findings to determine what they charge and cover for policyholders [86, 88].

One significant concern is the potential for discrimination [75, 86], something interviewees also raised. As one explained, “I would have a problem with [the insurance company] trying to discriminate [against] me based on how much sleep I get. Or, like, how much running I do.” Interviewees recognized that PFI flows to insurance companies could leave subjects worse off. As another explained: “I don’t trust insurance companies to make decisions that are based on my best interests. I think the more data they have about me, the more concerned they’d be about insuring me.” Another recognized that while they might personally benefit from such information flows, others—including those already enduring marginalization—may not. They explained: “I’m probably going to be on the beneficiary side [of PFI flows to insurance companies], so in principle, I should be for it, but I know that it’ll also really disadvantage a lot of people. People will get penalized. ... So, as a sort of civil rights issue, I’m going to step back and say even though I might benefit from it, overall it’s probably not a good trend.”

Indeed, the pricing structures of insurance policies that link premiums with behavioral tracking tend to reflect income rather than risk, such that wealthier policyholders end up getting discounted rates while poorer ones pay more [86]. The notion of essential services like health care costing more for those who can least afford them “will strike many as being inconsistent with our notions of economic and social justice” [13:146].

Step 8 (Evaluation 2): Consider the meaning or significance of moral and political factors in light of contextual values, ends, purposes, and goals.

This step asks researchers to assess how the practice under study could affect the context. Where Step 7 focused on what is at stake with the practice more generally, Step 8 situates those implications within the context where the practice occurs. When practices uphold the aims or ends of a context, they contribute to its integrity. In other words, they keep the context together. Conversely, when practices undermine or contradict the aims or ends of a context, they weaken its integrity. Thus, Step 8 is an assessment of contextual integrity.

Researchers can conduct this evaluation by considering the following questions: How does the practice under study align with the aims or goals of the context? How do the moral/political implications of the practice under study advance or undermine these contextual aims? In effect, this step involves integrating the findings from Step 5 (contextual norms) and Step 7 (moral/political implications of the practice) with information from Step 2 (the aims and goals of the context) to determine the extent to which the practice may be considered contextually appropriate. As with the other steps in the prescriptive portion of the CI framework, researchers may find it valuable to interpret their data using information from external sources, such as laws, scholarship, media coverage, corporate documents, or other cultural references that may shape a context.

Step 8 Example: The values of the health care context include “alleviating physical suffering, curing illness, and promoting the health of individuals as well as collectives” [64:134]. However, the way these values play out in practice is highly contested: “[W]e may disagree over whether prevention is more important than cure, prolonging individual life more important than average population health, and so forth. Some have argued that healthcare values include equity, the provision of care (or organs for transplants) according

to need, irrespective of ability to pay; others disagree. Some hold that physicians should respect whatever paths patients choose; others insist that physicians have a right and duty to steer” [66:227]. Thus, one avenue for assessing how automated transmission of PFI flows to doctors and insurance companies could affect contextual integrity is to consider the ways such flows align with these different values.

As a company, Fitbit has positioned wearable fitness tracking as supporting preventative and diagnostic measures [85]. On the preventative side, Fitbit has partnered with insurance companies [23] as well as Medicare and Medicaid, two U.S.-government-run programs that primarily insure older adults and people with low incomes, respectively, which the company frames as its effort to engage underserved populations [85]. Its Fitbit Care platform is designed to help organizations like employers and health plans [21] “understand their [members’] health behaviors at a deeper level, beyond what EHR [electronic health records] and claims data alone can provide” [22]. Fitbit created the program because “[p]ayers, providers and consumers alike know that digital tools to support and understand daily health behaviors are key to driving down rates of disease and costs” [22].

Fitbit’s comments frame the transmission of PFI to insurance companies as a way to foster the kinds of behaviors that support health, but also provide financial benefit. A few interviewees subscribed to this logic; those who took pride in their “good step numbers” and “active lifestyle” felt PFI flows to insurance companies could demonstrate that “I’m doing my preventative care.” Others wanted to see some of that financial benefit for themselves. For example, one said, “I think [PFI flows to insurance companies] would have to serve a purpose. I know that at least the state insurances, they do the annual exam and the blood test and stuff for a \$250 gift card or something like that.” It is worth noting that this desire for a benefit in exchange for data is not necessarily connected to a health-oriented purpose; this interviewee would support PFI transmission if they gained something, not because it could improve their health.

While incentivizing people to engage in healthy behaviors may not be inherently problematic, Step 7 explained why doing so in the context of health insurance can raise moral concerns [48]. If such practices make it harder for people to pay for health coverage and get treatment for health conditions, they could undermine the value of care in the healthcare context, and thus violate contextual integrity. In this vein, one interviewee wondered about the way that PFI flows to insurance companies could implicate doctors: “I’m not sure what the conditions are for treating me and obtaining coverage through my insurance. [The doctor] may be beholden to the insurance company in some way that I may not be aware of.” Indeed, Fitbit has worked with Google Cloud’s Healthcare API to promote interoperability with healthcare providers [62], and a Fitbit executive said the company seeks more integration with healthcare providers [85]. Despite issues about the accuracy of fitness tracker data, as noted in Step. 6, PFI can be used to make inferences about a subject’s personal relationships [89], diet, stress, and addictive behaviors [76]; and chronic conditions such as diabetes, high cholesterol, high blood pressure, or sleep apnea [9]. Thus, questions about how PFI flows could affect treatment itself are also important to consider.

On the diagnostic side, Fitbit has partnered with pharmaceutical companies and medical researchers to study and promote the use of its devices to detect atrial fibrillation (a-fib), an irregular heartbeat condition [25, 49, 50, 85]. In this case, Fitbit data is not automatically transmitted to doctors; rather, Fitbit notifies users in whom it detects a potential sign of a-

fib and encourages them to contact their doctor. This aligns with one interviewee's vision of an appropriate flow of PFI to doctors; they expressed a willingness to share PFI with their doctor when addressing a specific medical concern, "so then we could troubleshoot and put together a plan if I needed to change lifestyle decisions." It also appears to reflect the norm discussed in Step 5, in that a subject may choose to meet with a doctor and discuss their heart rate data, rather than have the heart rate data automatically sent to a doctor. If the threat of a-fib was extraordinary enough, CI contends that Fitbit *could* be justified in changing the transmission principle and transmitting the data to doctors automatically. But considering the value interviewees placed on being able to contextualize PFI data with their doctor, this may not be the case.

This analysis aligns with Step 6's concern and suggests that automated transmission of PFI for the purpose of providing health care could threaten contextual integrity by orienting subjects toward market-based incentives rather than health care benefits and could threaten their sense of autonomy in the doctor-patient relationship.

Step 9: Recommendation for or against the system or practice being studied.

The final step in Nissenbaum's decision heuristic asks researchers to judge, based on their analysis, what the course of action should be. CI acknowledges that privacy is not absolute. A new technology-driven practice may violate privacy, but if the practice advances a societal goal that outweighs the harm, it may be acceptable. For instance, the flow of health data to public agencies for a public purpose may raise concerns, but if a public health crisis like the COVID-19 pandemic arises, the societal interest in mitigating the spread of an infectious disease may outweigh such concerns and thus justify the information flows [28, 93].

Researchers can offer a recommendation by considering the following questions: Should the practice under study continue as is, or should it be rejected? What, if any, modifications or conditions should be implemented? Carrying the CI analysis through to this recommendation step is important so that privacy doesn't become what Nissenbaum calls a "flamboyant smokescreen" against potentially worthwhile information flows [3]. She used such terms when Apple and Google invoked privacy concerns to explain their resistance toward working with public health agencies on coronavirus tracking, noting the irony that both companies tolerate surveillance in other capacities. Indeed, while Apple and Google rightfully implemented individual privacy protections into their tracing technologies, their privacy concerns about health agency involvement were perhaps less warranted given the important and trusted position health agencies play in contact tracing [34]. These nuances demonstrate why privacy needs to be put, as Nissenbaum [64] notes, in context.

Step 9 Example: This example analysis examined the flow of PFI in the health care context. It established that the flow of patients (*senders*) choosing (*transmission principle*) to provide PFI (*information type*) about themselves (*subjects*) to doctors (*recipients*) who need it (*transmission principle*) as a norm of health care and observed that the automated transmission of PFI could pose concerns. In particular, the increased integration of automated PFI flows in the larger healthcare ecosystem can introduce new *senders* (Step 3) and *transmission principles* (Step 4-5) into the healthcare context, provide more granularity to information actors may already receive (Step 5), raise questions about diagnosis and

treatment, given the unreliability of PFI (Step 6), and unfairly limit people who already experience marginalization from obtaining affordable or complete medical coverage (Step 7), potentially undermining the aims of health care itself (Step 8).

As companies like Fitbit pivot their focus from consumer electronics to digital health [87], their practices deserve scrutiny from a contextual integrity perspective. Complicating things further is Google's acquisition of Fitbit [77] which provides the internet giant access to PFI from millions of users. While Google has promised not to use Fitbit data for targeted advertising [29], questions remain how else Google might integrate PFI across other products, including a possible move to enter the healthcare space itself [6]. Such possibilities add fuel to concerns about the automated flow of PFI for purposes beyond those a subject expects when they begin using a fitness tracker.

Our example CI analysis indicates that there are people who find the intersection between PFI and health promising. However, the analysis in Steps 6-8 suggest that PFI should not automatically flow to insurance companies, given its potential for exacerbating inequality, and that the flow of PFI into health care is something patients should individually pursue with their doctors, rather than something that doctors should mandate of patients. This approach strikes a balance between maintaining the integrity of the healthcare context while permitting those who value fitness tracking to integrate it into healthcare on their terms.

4 BEYOND THE ROADMAP: ADDITIONAL CONCERNS WHEN APPLYING CI

Having walked through the roadmap and example, in this section we explain how the roadmap addresses critiques of CI, identify opportunities for theory development to further strengthen CI, and offer suggestions for future work.

4.1 Addressing Critiques of CI

As CI has gained popularity in HCI and social computing research, scholars have critiqued the theory for insufficiently accounting for obscurity in institutional or automated practices [41, 84] as well as societal inequities [60]. Nissenbaum [66] has acknowledged CI's limitations, and certainly, no theory is perfect. However, we believe that some of these concerns can be addressed by engaging with the full CI framework, and we offer the roadmap presented above as a way to help researchers—particularly those conducting qualitative work—do so.

In her study of the privacy implications of direct-to-consumer genetic testing, King explains that CI “is limited in explaining why individuals disclose personal information in digital contexts where the norms are emergent, in flux, or when a company's practices may violate contextual norms” [41:6]. She adds that “the fact that the respondents had not as of yet experienced a privacy violation by the company as conceptualized by contextual integrity does not mean that no privacy risks were present; it is that many of the privacy risks are inherent to the state of individual data being connected to others' data” [41:25]. Similarly, Skeba and Baumer suggest that the social concerns of police use of facial recognition technologies, such as “potential stifling of free speech and assembly, disparate negative impacts on minority communities, and arrests made based on false-positive matches...cannot easily be articulated in the language of CI, though, since [they] are based more on perceptions about possible actions than on any discrete act that might constitute a violation” [84:7]. However, Skeba and Baumer's [84] explanation of the CI framework focuses on parameters and norms and makes no mention of the framework's second half, in which

norm violations are evaluated against moral and political values. Indeed, applying the full, nine-step CI heuristic to direct-to-consumer genetic testing and facial recognition practices would likely identify the concerns these scholars have referenced. CI does not require an individual to experience a privacy violation in order for the information flow to be deemed concerning. Steps 1-5 provide a structure for describing a given information flow, and Steps 6-9 involve comparing the information flow against moral and political values, which goes beyond individuals' experiences to consider the wider landscape, including legal and scholarly perspectives. Flows that contravene social values (regardless of individual preferences or experiences) can be judged problematic and resisted.

Of course, the question of whose values and judgments are heeded is important. This concern grounds McDonald and Forte's [60] critique of CI. They explain that "[n]orms are useful shorthand for how people expect to navigate relationships, situations, and spaces but...it is often the most privileged individuals who are able to participate in norm-setting and articulation" [60:3]. Indeed, long-standing structural inequities mean that many societies privilege the worldviews of certain groups over others, something CI does not directly address. McDonald and Forte contend that "[c]ontextual integrity assumes that in a given context the impacts of privacy violations will be experienced equally" [60:9]. While it is true that CI's parameters do not address dimensions linked to structural inequality (e.g., race, gender, class, etc.), the second half of the framework is meant to help analysts uncover differential impacts. In our example analysis evaluating the flow of PFI in the healthcare context, Step 7 centered on concerns of discrimination—which critical scholars like Khiara Bridges [13] unpack in their evaluations of the link between privacy rights and socioeconomic status. Harnessing the full CI framework and complementing the prescriptive analysis with insights from critical theories like intersectionality, as McDonald and Forte [60] advocate, can bring privacy research closer to Nissenbaum's goal of having CI help "reveal the delicate balance of multiple conflicting interests and plurality of values that complex constraints on flow seek to realize" [66:234].

4.2 Opportunities for Theory Refinement and Expansion

While harnessing the full framework, as we have advocated in this paper, will go a long way toward addressing critiques of CI, there remain several opportunities to further strengthen the theory. We focus on two areas where we think there are opportunities for refining Nissenbaum's framework: parameters and the connections between parameters, flows, and norms.

4.2.1 Open Questions About CI's Parameters. The five CI parameters—actors (further divided into subject, sender, and recipient), information types, and transmission principles—are the building blocks of information flows. But open questions remain about what the parameters cover and how they are connected. Nissenbaum [64] acknowledges that actors can be individuals, groups, or collective entities like companies. But what about artifacts like Fitbits or technological systems like TikTok's "For You" algorithm? As social practices and digital technologies become more deeply entangled, it seems only logical that technologies can stand as actors. Relatedly, the medium of information transmission (e.g., face-to-face conversation, handwritten note, text message, public social media post) has privacy implications, but it's not clear whether or how the CI parameters address this. Some have linked medium with the transmission principle parameter [46], while others suggest it may need to be a separate parameter [44]. Better integrating non-human, technological actors into CI could also help push privacy research beyond a focus on individual preferences, expectations, and controls and toward more system-driven understandings. The CSCW community, with its commitment to theories such as structuration and sociomateriality [69, 70], is well-positioned for such work.

As we note in Section 3, the transmission principle parameter is among CI's most valuable elements, but also one of its most misunderstood. In our roadmap, we aim to provide some conceptual clarity in Step 4 by encouraging researchers to approach the transmission principle as the sentiment that follows the expression, "This information flow is fine IF..." We believe this approach is compatible with those like Shvartzshnaider et al.'s [82], which divides the transmission principle into the sub-parameters of aims, conditions, modalities, and consequences. However, our narrow focus on the conditions of appropriateness contrasts with Frik et al.'s [27] broader interpretation of transmission principles as encompassing the purposes, benefits, and risks of information disclosure. We believe that focusing on conditions of appropriateness more closely matches the definition of a transmission principle as a "constraint on the flow (distribution, dissemination, transmission) of information from party to party in a context" [64:145], and that purposes, benefits, and risks fit better in other steps of the framework. Nissenbaum [66] aligns purposes with context and as such, we incorporated it into Step 2. Benefits and risks become important when considering whether the flow as a whole is worthwhile, and thus we embedded them in Steps 7 and 8. However, we echo Frik et al. [27] in calling for more theory development surrounding transmission principles. For instance, what happens when multiple transmission principles govern an information flow? How do changes in actor or information type affect transmission principles? This speaks to the need for more clarity about how the pieces of CI, from parameters to flows to norms, are themselves connected. We address this in the next section.

4.2.2 Conceptual Clarity Regarding Connections Between Parameters, Flows, Norms, and Contexts. CI's parameters (actors, information types, and transmission principles) are the ingredients of information flows. CI contends that when information flows align with contextual norms, the social practice in question maintains contextual integrity. However, if flows do not align with contextual norms, privacy concerns may arise. Analysts can use CI's parameters to pinpoint where the misalignment occurs. To determine whether the information flow should be altered as a result of the discrepancy—or, if the privacy concern is significant enough to warrant action—analysts evaluate the misalignment against broader political, moral, and contextual values and offer a recommendation about how to act.

While this sequence may appear straightforward, applying it can be challenging. For instance, one application of CI discusses violations of parameters [84], but CI addresses violations of norms, not parameters. We hope that the roadmap and example case presented in this paper help demystify the different pieces of CI, but more work is needed to theorize the connections between them. For instance, at what point does something become a distinct information flow? In our example analysis, disclosing general fitness information in conversation with a doctor, showing Fitbit data to a doctor during an appointment, and transmitting Fitbit data to a medical file represent different information flows, and we can use CI's parameters to identify what differences exist between them. But when do changes in parameters create different information flows? Compare a patient voluntarily disclosing fitness information to a doctor and a doctor asking the patient for fitness information. Are these the same information flows with different transmission principles, or are they different information flows precisely because their transmission principles differ?

The connections between information flows and norms also merit further exploration. Is each information flow expected to have a correlating norm? Should analysts describe norms with the same granularity as information flows? Since norms are contextually driven, how should analysts handle situations where social practices engage overlapping contexts? This question is

particularly important considering that context collapse, or the merging and blurring of multiple contexts in a given sphere of interaction, is a defining condition of the contemporary digital environment [57]. The HCI and social computing community has taken up such questions [16, 51, 92], and we encourage continued exploration to develop this aspect of CI further.

Finally, the connections between preferences, expectations, and norms must be clarified. Critiques of CI suggest the framework focuses too much on individual preferences [41, 84], but Nissenbaum [64, 66] emphasizes that norms reflect fundamentally collective, not individual, expectations. One source for the confusion may be that most analytical applications of CI have been statistical. Factorial vignette surveys that operationalize CI ask respondents to rate a series of scenarios for appropriateness [e.g., 54, 91]. One would expect a certain level of variation at the individual level, as different people understandably have different interpretations of appropriateness. But in a well-designed survey, overarching population norms should appear, smoothing out individual variations (unless the scenario is truly so novel that no entrenched norms exist). Furthermore, researchers can visualize their survey data to pinpoint the most salient parameters and explain the extent to which they influence judgments [4, 5]. Qualitative research paradigms are not intended to generalize data from a sample to a population and thus will employ different methods for identifying and studying norms. One of the strengths of qualitative research is its ability to interpret how meaning and experience shape a given topic of study. Our roadmap suggests that researchers identify the CI parameters in their data and use information from various domains, including law, policy, culture, scholarship, media coverage, and corporate documents to make sense of the preferences, expectations, and norms relevant to their practice under study. But we encourage researchers and methodologists to continue developing guidance to support qualitative researchers in this work.

4.3 Future Work

We invite scholars to build on the roadmap presented in this paper and continue examining the role of CI in HCI and social computing research. Since qualitative research encompasses a range of epistemological commitments and data analysis methods, we tried to avoid making our roadmap overly prescriptive. However, since coding is an integral part of many qualitative methodologies, future work could offer more detailed guidance about how to code qualitative data for information relevant to each step of the CI framework. In light of section 5.2.2's discussion of the need for more clarity about the links and distinctions between information parameters, flows, and norms, this guidance will need to consider how to reconcile the reductive nature of coding with the fluid and dynamic nature of information flows. The research community would also benefit from reflections on how to incorporate CI into non-coding-based methodologies like ethnography or case studies.

To gain a deeper understanding of CI's role in HCI and social computing research, we also encourage researchers to conduct a systematic review of CI in this space, similar to Benthall et al.'s [11] review of CI-based studies in computer science. Badillo-Urquiola et al.'s [7] preliminary analysis offers a starting point, but a full review is needed to better discern CI's utility—and limitations—as well as identify opportunities for theoretical and methodological development.

5 CONCLUSION

As an analytical framework, contextual integrity enables researchers to pinpoint how information flows raise privacy concerns. But its power extends beyond such descriptive precision, as a CI analysis also lays the groundwork for helping people determine how to respond

to such concerns. Indeed, the idea that digital technologies raise privacy concerns borders on cliché. Privacy researchers, with their expert knowledge and methodological skills, are well-positioned to help society figure out what to do about those concerns. Conducting a full CI analysis using Nissenbaum's nine-step decision heuristic can equip researchers to make specific recommendations about how to alter a technology-driven practice to make it more privacy preserving—or perhaps advocate that the system or practice should instead be resisted outright. We hope that the roadmap and example analysis provided in this paper help privacy researchers harness the full potential of CI.

ACKNOWLEDGMENTS

A portion of this work was presented at the 2021 Symposium on Applications of Contextual Integrity, and we thank the audience for their useful feedback. We thank Janet Ruppert for contributing to our many conversations about CI, and we thank Yuting Liao and Katie Chamberlain Kritikos for their assistance conducting interviews. We are grateful to Heather Patterson for generously sharing her data and championing this work.

REFERENCES

- [1] Noura Abdi, Xiao Zhan, Kopo M. Ramokapane, and Jose Such. 2021. Privacy Norms for Smart Home Personal Assistants. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, May 06, 2021. ACM, Yokohama Japan, 1–14. <https://doi.org/10.1145/3411764.3445122>
- [2] Imtiaz Ahmad, Rosta Farzan, Apu Kapadia, and Adam J. Lee. 2020. Tangible Privacy: Towards User-Centric Sensor Designs for Bystander Privacy. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2 (October 2020), 1–28. <https://doi.org/10.1145/3415187>
- [3] Reed Albergotti and Drew Harwell. 2020. Apple and Google are building a virus-tracking system. Health officials say it will be practically useless. *Washington Post*. Retrieved January 6, 2023 from <https://www.washingtonpost.com/technology/2020/05/15/app-apple-google-virus/>
- [4] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering Smart Home Internet of Things Privacy Norms Using Contextual Integrity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (June 2018), 1–23. <https://doi.org/10.1145/3214262>
- [5] Noah Apthorpe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus {COPPA}. In *28th USENIX Security Symposium*, 2019. 123–140. Retrieved from <https://www.usenix.org/conference/usenixsecurity19/presentation/apthorpe>
- [6] Patrick Lucas Austin. 2019. The Real Reason Google Is Buying Fitbit. *Time*. Retrieved October 14, 2023 from <https://time.com/5717726/google-fitbit/>
- [7] Karla Badillo-Urquiola, Xinru Page, and Pamela Wisniewski. 2018. Literature Review: Examining Contextual Integrity within Human-Computer Interaction. September 2018. Princeton, N.J., 1–4.
- [8] Karla Badillo-Urquiola, Yaxing Yao, Oshrat Ayalon, Bart Knijnenburg, Xinru Page, Eran Toch, Yang Wang, and Pamela J. Wisniewski. 2018. Privacy in Context: Critically Engaging with Theory to Guide Privacy Research and Design. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing*, October 30, 2018. ACM, Jersey City NJ USA, 425–431. <https://doi.org/10.1145/3272973.3273012>
- [9] Brandon Ballinger, Johnson Hsieh, Avesh Singh, Nimit Sohoni, Jack Wang, Geoffrey H. Tison, Gregory M. Marcus, Jose M. Sanchez, Carol Maguire, Jeffrey E. Olgin, and Mark J. Pletcher. 2018. DeepHeart: Semi-supervised sequence learning for cardiovascular risk prediction. In *AAAI 2018*. Retrieved April 19, 2018 from <https://arxiv.org/abs/1802.02511>
- [10] Louise Barkhuus. 2012. The Mismeasurement of Privacy: Using Contextual Integrity to Reconsider Privacy in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*, 2012. ACM, New York, NY, USA, 367–376. <https://doi.org/10.1145/2207676.2207727>
- [11] Sebastian Benthall, Seda Gürses, and Helen Nissenbaum. 2017. Contextual Integrity through the Lens of Computer Science. *Foundations and Trends® in Privacy and Security* 2, 1 (2017), 1–69. <https://doi.org/10.1561/33000000016>
- [12] Anne Bowser, Katie Shilton, Jenny Preece, and Elizabeth Warrick. 2017. Accounting for Privacy in Citizen Science: Ethical Research in a Context of Openness. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*, 2017. ACM, New York, NY, USA, 2124–2136. <https://doi.org/10.1145/2998181.2998305>
- [13] Khiara M. Bridges. 2017. *The Poverty of Privacy Rights*. Stanford University Press.

- [14] Kristen V. Brown. 2017. When an insurer sells you an Apple Watch For \$25, how much are you giving away? *Gizmodo*. Retrieved April 19, 2018 from <https://gizmodo.com/when-an-insurer-sells-you-an-apple-watch-for-25-how-m-1819912738>
- [15] Yunan Chen and Heng Xu. 2013. Privacy management in dynamic groups: understanding information privacy in medical practices. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, 2013. ACM Press, 541–552. <https://doi.org/10.1145/2441776.2441837>
- [16] Hichang Cho and Anna Filippova. 2016. Networked Privacy Management in Facebook: A Mixed-Methods and Multinational Study. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, February 27, 2016. ACM, San Francisco California USA, 503–514. <https://doi.org/10.1145/2818048.2819996>
- [17] Michelle M. Christovich. 2016. Why Should We Care What Fitbit Shares - A Proposed Statutory Solution to Protect Sensitive Personal Fitness Information. *Hastings Comm. & Ent. L.J.* 38, (2016), 91–116.
- [18] Melanie Duckert and Louise Barkhuus. 2022. Protecting Personal Health Data through Privacy Awareness: A study of perceived data privacy among people with chronic or long-term illness. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP (January 2022), 1–22. <https://doi.org/10.1145/3492830>
- [19] Lynne M Feehan, Jasmina Geldman, Eric C Sayre, Chance Park, Allison M Ezzat, Ju Young Yoo, Clayon B Hamilton, and Linda C Li. 2018. Accuracy of Fitbit Devices: Systematic Review and Narrative Syntheses of Quantitative Data. *JMIR Mhealth Uhealth* 6, 8 (August 2018), e10527. <https://doi.org/10.2196/10527>
- [20] Fitbit. How do I connect Fitbit with another app? *Fitbit Help*. Retrieved October 14, 2023 from https://help.fitbit.com/articles/en_US/Help_article/1742.htm
- [21] Fitbit. What should I know about participating in a Fitbit Care program? *Fitbit Help*. Retrieved October 14, 2023 from https://help.fitbit.com/articles/en_US/Help_article/1760.htm
- [22] Fitbit. Fitbit’s Population Health and Wellness Solution: Fitbit Care. *Fitbit Health Solutions*. Retrieved October 14, 2023 from <https://healthsolutions.fitbit.com/what-is-fitbit-care/>
- [23] Fitbit Team. 2016. Fitbit introduces “Fitbit Group Health.” *The Keyword*. Retrieved October 14, 2023 from <https://blog.google/products/fitbit/fitbit-introduces-fitbit-group-health/>
- [24] Fitbit Team. 2018. Fitbit and Google partner to innovate digital health, wearables. *The Keyword*. Retrieved October 14, 2023 from <https://blog.google/products/fitbit/fitbit-and-google-partner-innovate-digital-health-wearables/>
- [25] Fitbit Team. 2019. Bristol-Myers Squibb-Pfizer Alliance and Fitbit collaborate. *The Keyword*. Retrieved October 14, 2023 from <https://blog.google/products/fitbit/bristol-myers-squibb-pfizer-alliance-and-fitbit-collaborate/>
- [26] Susannah Fox and Maeve Duggan. 2013. Tracking for health. *Pew Research Center: Internet & Technology*. Retrieved November 21, 2016 from <http://www.pewinternet.org/2013/01/28/tracking-for-health/>
- [27] Alisa Frik, Julia Bernd, and Serge Egelman. 2022. A Model of Contextual Factors Affecting Older Adults’ Information-Sharing Decisions in the US. *ACM Trans. Comput.-Hum. Interact.* (August 2022), 3557888. <https://doi.org/10.1145/3557888>
- [28] Frederic Gerdon, Helen Nissenbaum, Ruben L. Bach, Frauke Kreuter, and Stefan Zins. 2021. Individual Acceptance of Using Health Data for Private and Public Benefit: Changes During the COVID-19 Pandemic. *Harvard Data Science Review* (April 2021). <https://doi.org/10.1162/99608f92.edf2fc97>
- [29] Shirin Ghaffary and Rani Molla. 2019. Google says it won’t use your Fitbit data to target you with ads. But what else will it do? *Vox*. Retrieved October 14, 2023 from <https://www.vox.com/recode/2019/11/1/20943583/google-fitbit-acquisition-privacy-antitrust>
- [30] Sarah Gilbert, Jessica Vitak, and Katie Shilton. 2021. Measuring Americans’ Comfort With Research Uses of Their Social Media Data. *Social Media + Society* 7, 3 (July 2021), 205630512110338. <https://doi.org/10.1177/20563051211033824>
- [31] Liang Gou, Michelle X. Zhou, and Huahai Yang. 2014. KnowMe and ShareMe: understanding automatically discovered personality traits from social media and user sharing preferences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, April 26, 2014. ACM, Toronto Ontario Canada, 955–964. <https://doi.org/10.1145/2556288.2557398>
- [32] Audrey Guinchard. 2018. Taking proportionality seriously: The use of contextual integrity for a more informed and transparent analysis in EU data protection law. *European Law Journal* 24, 6 (November 2018), 434–457. <https://doi.org/10.1111/eulj.12273>
- [33] Julia Hanson, Miranda Wei, Sophie Veys, Matthew Kugler, Lior Strahilevitz, and Blase Ur. 2020. Taking Data Out of Context to Hyper-Personalize Ads: Crowdworkers’ Privacy Perceptions and Decisions to Disclose Private Information. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, April 21, 2020. ACM, Honolulu HI USA, 1–13. <https://doi.org/10.1145/3313831.3376415>
- [34] Eszter Hargittai, Elissa M. Redmiles, Jessica Vitak, and Michael Zimmer. 2020. Americans’ willingness to adopt a COVID-19 tracking app: The role of app distributor. *First Monday* (October 2020). <https://doi.org/10.5210/fm.v25i11.11095>

- [35] Jane Henriksen-Bulmer. 2019. Incorporating contextual integrity into privacy decision making: a risk based approach. Doctoral Thesis. Bournemouth University. Retrieved January 8, 2024 from <https://eprints.bournemouth.ac.uk/32385/>
- [36] Gordon Hull, Heather Richter Lipford, and Celine Latulipe. 2011. Contextual gaps: privacy issues on Facebook. *Ethics Inf Technol* 13, 4 (December 2011), 289–302. <https://doi.org/10.1007/s10676-010-9224-8>
- [37] Jack Jamieson, Naomi Yamashita, Daniel A. Epstein, and Yunan Chen. 2021. Deciding If and How to Use a COVID-19 Contact Tracing App: Influences of Social Factors on Individual Use in Japan. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (October 2021), 1–30. <https://doi.org/10.1145/3479868>
- [38] Kyle M. L. Jones and Ellen LeClere. 2018. Contextual expectations and emerging informational harms: A primer on academic library participation in learning analytics initiatives. In *Applying library values to emerging technology: decision-making in the age of open access, maker spaces, and the ever-changing library*, Peter D. Fernandez and Kelly Tilton (eds.). Association of College and Research Libraries, a division of the American Library Association, Chicago, Illinois, 357–371.
- [39] Katherine Keisler-Starkey and Lisa N. Bunch. 2021. *Health Insurance Coverage in the United States: 2020*. U.S. Census Bureau, Washington, DC. Retrieved from <https://www.census.gov/library/publications/2021/demo/p60-274.html>
- [40] Paula Kift and Helen Nissenbaum. 2016. Metadata in Context - An Ontological and Normative Analysis of the NSA's Bulk Telephony Metadata Collection Program. *ISJLP* 13, 2 (2017 2016), 333–372.
- [41] Jennifer King. 2019. "Becoming Part of Something Bigger": Direct to Consumer Genetic Testing, Privacy, and Personal Disclosure. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (November 2019), 1–33. <https://doi.org/10.1145/3359260>
- [42] Bart P. Knijnenburg, Xinru Page, Pamela Wisniewski, Heather Richter Lipford, Nicholas Proferes, and Jennifer Romano (Eds.). 2021. *Modern socio-technical perspectives on Privacy*. Springer Nature, Gewerbestrasse, Switzerland.
- [43] Gary L. Kreps. 1988. The Pervasive Role of Information in Health and Health Care: Implications for Health Communication Policy. *Annals of the International Communication Association* 11, 1 (January 1988), 238–276. <https://doi.org/10.1080/23808985.1988.11678690>
- [44] Priya C. Kumar and Virginia L. Byrne. 2022. The 5Ds of privacy literacy: a framework for privacy education. *ILS* 123, 7/8 (August 2022), 445–461. <https://doi.org/10.1108/ILS-02-2022-0022>
- [45] Priya C. Kumar, Mega Subramaniam, Jessica Vitak, Tamara L. Clegg, and Marshini Chetty. 2020. Strengthening Children's Privacy Literacy through Contextual Integrity. *Media and Communication* 8, 4 (November 2020), 175–184. <https://doi.org/10.17645/mac.v8i4.3236>
- [46] Priya Kumar, Shalmali Milind Naik, Utkarsha Ramesh Devkar, Marshini Chetty, Tamara L. Clegg, and Jessica Vitak. 2017. "No telling passcodes out because they're private": Understanding children's mental models of privacy and security online. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 64:1-64:21. <https://doi.org/10.1145/3134699>
- [47] Heather Richter Lipford, Gordon Hull, Celine Latulipe, Andrew Besmer, and Jason Watson. 2009. Visible Flows: Contextual Integrity and the Design of Privacy Mechanisms on Social Network Sites. 2009. IEEE, 985–989. <https://doi.org/10.1109/CSE.2009.241>
- [48] Michele Loi, Christian Hauser, and Markus Christen. 2022. Highway to (Digital) Surveillance: When Are Clients Coerced to Share Their Data with Insurers? *J Bus Ethics* 175, 1 (January 2022), 7–19. <https://doi.org/10.1007/s10551-020-04668-1>
- [49] Steven A. Lubitz, Anthony Z. Faranesh, Steven J. Atlas, David D. McManus, Daniel E. Singer, Sherry Pagoto, Alexandros Pantelopoulous, and Andrea S. Foulkes. 2021. Rationale and design of a large population study to validate software for the assessment of atrial fibrillation from data acquired by a consumer tracker or smartwatch: The Fitbit heart study. *American Heart Journal* 238, (August 2021), 16–26. <https://doi.org/10.1016/j.ahj.2021.04.003>
- [50] Steven A. Lubitz, Anthony Z. Faranesh, Caitlin Selvaggi, Steven J. Atlas, David D. McManus, Daniel E. Singer, Sherry Pagoto, Michael V. McConnell, Alexandros Pantelopoulous, and Andrea S. Foulkes. 2022. Detection of Atrial Fibrillation in a Large Population Using Wearable Devices: The Fitbit Heart Study. *Circulation* 146, 19 (November 2022), 1415–1424. <https://doi.org/10.1161/CIRCULATIONAHA.122.060291>
- [51] Ameera Mansour and Helena Francke. 2021. Collective Privacy Management Practices: A Study of Privacy Strategies and Risks in a Private Facebook Group. *Proc. ACM Hum.-Comput. Interact.* 5, (October 2021), 1–27. <https://doi.org/10.1145/3479504>
- [52] Stephen T. Margulis. 2011. Three Theories of Privacy: An Overview. In *Privacy Online*, Sabine Treppe and Leonard Reinecke (eds.). Springer Berlin Heidelberg, 9–17. Retrieved April 6, 2014 from http://link.springer.com.proxy.lib.umich.edu/chapter/10.1007/978-3-642-21521-6_2
- [53] Kirsten E. Martin. 2012. Diminished or just different?: A factorial vignette study of privacy as a social contract. *Journal of Business Ethics* 111, 4 (2012), 519–539.

- [54] Kirsten Martin and Helen Nissenbaum. 2016. Measuring Privacy: An Empirical Test Using Context to Expose Confounding Variables. *Colum. Sci. & Tech. L. Rev.* 18, (2017 2016), 176–218.
- [55] Kirsten Martin and Helen Nissenbaum. 2017. Privacy Interests in Public Records: An Empirical Investigation. *Harvard Journal of Law and Technology* 31, 1 (2017), 111–143.
- [56] Kirsten Martin and Helen Nissenbaum. 2020. What is it about Location? *Berkeley Technology Law Journal* 35, 1 (2020), 252–326.
- [57] Alice E. Marwick and danah boyd. 2011. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media Society* 13, 1 (February 2011), 114–133. <https://doi.org/10.1177/1461444810365313>
- [58] Alice E Marwick and danah boyd. 2014. Networked privacy: How teenagers negotiate context in social media. *New Media & Society* 16, 7 (November 2014), 1051–1067. <https://doi.org/10.1177/1461444814543995>
- [59] Alice Emily Marwick. 2023. *The private is political: networked privacy and social media*. Yale University Press, New Haven.
- [60] Nora McDonald and Andrea Forte. 2020. The Politics of Privacy Theories: Moving from Norms to Vulnerabilities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, April 21, 2020. ACM, Honolulu HI USA, 1–14. <https://doi.org/10.1145/3313831.3376167>
- [61] Darakhshan J. Mir. 2021. Designing for the Privacy Commons. In *Governing Privacy in Knowledge Commons* (1st ed.), Madelyn Rose Sanfilippo, Brett M. Frischmann and Katherine J. Strandburg (eds.). Cambridge University Press, 245–267. <https://doi.org/10.1017/9781108749978.011>
- [62] Gregory J. Moore. 2018. New collaboration with Fitbit to drive positive health outcomes. *Google Cloud Blog*. Retrieved October 14, 2023 from <https://cloud.google.com/blog/topics/inside-google-cloud/new-collaboration-fitbit-drive-positive-health-outcomes>
- [63] National Institutes of Health. 2023. Research Roundup: All of Us Participants’ Fitbit Data Drive New Research. *All of Us Research Program | NIH*. Retrieved October 14, 2023 from <https://allofus.nih.gov/news-events/announcements/research-roundup-all-us-participants-fitbit-data-drive-new-research>
- [64] Helen Nissenbaum. 2010. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, Stanford, CA.
- [65] Helen Nissenbaum. 2015. Respect for context as a benchmark for privacy online: what it is and isn’t. In *Social Dimensions of Privacy* (1st ed.), Beate Roessler and Dorota Mokrosinska (eds.). Cambridge University Press, 278–302. <https://doi.org/10.1017/CBO9781107280557.016>
- [66] Helen Nissenbaum. 2019. Contextual integrity up and down the data food chain. *Theoretical Inquiries in Law* 20, 1 (March 2019), 221–256. <https://doi.org/10.1515/til-2019-0008>
- [67] Aletta Norval and Elpida Prasopoulou. 2017. Public faces? A critical exploration of the diffusion of face recognition technologies in online social networks. *New Media & Society* 19, 4 (April 2017), 637–654. <https://doi.org/10.1177/1461444816688896>
- [68] Leysan Nurgalieva, Seamus Ryan, Andreas Balaskas, Janne Lindqvist, and Gavin Doherty. 2022. Public Views on Digital COVID-19 Certificates: a Mixed Methods User Study. In *CHI Conference on Human Factors in Computing Systems*, April 27, 2022. ACM, New Orleans LA USA, 1–28. <https://doi.org/10.1145/3491102.3502066>
- [69] Wanda J. Orlikowski. 2008. Using Technology and Constituting Structures: A Practice Lens for Studying Technology in Organizations. In *Resources, Co-Evolution and Artifacts*. Springer London, London, 255–305. https://doi.org/10.1007/978-1-84628-901-9_10
- [70] Carsten S. Østerlund, Pernille Bjørn, Paul Dourish, Richard Harper, and Daniela K. Rosner. 2015. Sociomateriality and Design. In *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing*, February 28, 2015. ACM, Vancouver BC Canada, 126–130. <https://doi.org/10.1145/2685553.2699336>
- [71] Leysia Palen and Paul Dourish. 2003. Unpacking “Privacy” for a Networked World. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’03)*, 2003. ACM, New York, NY, USA, 129–136. <https://doi.org/10.1145/642611.642635>
- [72] Mika Pantzar and Minna Ruckenstein. 2017. Living the metrics: Self-tracking and situated objectivity. *DIGITAL HEALTH* 3, (2017), 1–10. <https://doi.org/10.1177/2055207617712590>
- [73] Heather Patterson. 2013. Contextual expectations of privacy in self-generated health information flows. March 30, 2013. TPRC, Arlington, VA, 1–48. Retrieved October 31, 2017 from <http://dx.doi.org/10.2139/ssrn.2242144>
- [74] Heather Patterson and Helen Nissenbaum. 2013. Context-dependent expectations of privacy in self-generated mobile health data. June 06, 2013. .
- [75] Inbar Raber, Cian P. McCarthy, and Robert W. Yeh. 2019. Health Insurance and Mobile Health Devices: Opportunities and Concerns. *JAMA* 321, 18 (May 2019), 1767. <https://doi.org/10.1001/jama.2019.3353>
- [76] Andrew Raji, Animikh Ghosh, Santosh Kumar, and Mani Srivastava. 2011. Privacy Risks Emerging from the Adoption of Innocuous Wearable Sensors in the Mobile Environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’11)*, 2011. ACM, New York, NY, USA, 11–20. <https://doi.org/10.1145/1978942.1978945>

- [77] Akanksha Rana and Noor Zainab Hussain. 2019. Google taps fitness tracker market with \$2.1 billion bid for Fitbit. *Reuters*. Retrieved October 14, 2023 from <https://www.reuters.com/article/us-fitbit-m-a-alphabet-idUSKBN1XB47G>
- [78] Christopher Rowland. 2019. With fitness trackers in the workplace, bosses can monitor your every step — and possibly more. *Washington Post*. Retrieved January 14, 2023 from https://www.washingtonpost.com/business/economy/with-fitness-trackers-in-the-workplace-bosses-can-monitor-your-every-step--and-possibly-more/2019/02/15/75ee0848-2a45-11e9-b011-d8500644dc98_story.html
- [79] Madelyn Rose Sanfilippo, Brett M. Frischmann, and Katherine J. Strandburg (Eds.). 2021. *Governing Privacy in Knowledge Commons* (1st ed.). Cambridge University Press. <https://doi.org/10.1017/9781108749978>
- [80] Pan Shi, Heng Xu, and Yunan Chen. 2013. Using Contextual Integrity to Examine Interpersonal Information Boundary on Social Network Sites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 2013. ACM, New York, NY, USA, 35–38. <https://doi.org/10.1145/2470654.2470660>
- [81] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: perceptions of privacy and mobile app use. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, 2014. ACM Press, Toronto, Ontario, Canada, 2347–2356. <https://doi.org/10.1145/2556288.2557421>
- [82] Yan Shvartzshnaider, Madelyn Rose Sanfilippo, and Noah Apthorpe. 2022. GKC-CI: A unifying framework for contextual norms and information governance. *Asso for Info Science & Tech* 73, 9 (September 2022), 1297–1313. <https://doi.org/10.1002/asi.24633>
- [83] Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. 2016. Learning Privacy Expectations by Crowdsourcing Contextual Informational Norms. In *Fourth AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2016)*, 2016. AAAI, 209–218.
- [84] Patrick Skeba and Eric P. S. Baumer. 2020. Informational friction as a lens for studying algorithmic aspects of privacy. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2 (October 2020), 101:1–101:22. <https://doi.org/10.1145/3415172>
- [85] Gret Slabodkin. 2020. Fitbit looks to expand healthcare partnerships in 2020. *Health Data Management*. Retrieved October 14, 2023 from [https://www.healthdatamanagement.com/articles/\[object Object\]](https://www.healthdatamanagement.com/articles/[object Object])
- [86] Etye Steinberg. 2022. Run for Your Life: The Ethics of Behavioral Tracking in Insurance. *J Bus Ethics* 179, 3 (September 2022), 665–682. <https://doi.org/10.1007/s10551-021-04863-8>
- [87] Abigail Stevenson. 2016. Fitbit CEO reveals he's transforming the mission and purpose of the company. *CNBC*. Retrieved October 14, 2023 from <https://www.cnn.com/2016/10/06/fitbit-ceo-reveals-hes-transforming-the-mission-and-purpose-of-the-company.html>
- [88] Salvatore Tedesco, John Barton, and Brendan O'Flynn. 2017. A Review of Activity Trackers for Senior Citizens: Research Perspectives, Commercial Landscape and the Role of the Insurance Industry. *Sensors* 17, 6 (June 2017), 1277. <https://doi.org/10.3390/s17061277>
- [89] Kota Tsubouchi, Ryoma Kawajiri, and Masamichi Shimosaka. 2013. Working-relationship Detection from Fitbit Sensor Data. In *Proceedings of the 2013 ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication*, 2013. ACM, New York, NY, USA, 115–118. <https://doi.org/10.1145/2494091.2494123>
- [90] Christine Utz, Steffen Becker, Theodor Schnitzler, Florian M. Farke, Franziska Herbert, Leonie Schaewitz, Martin Degeling, and Markus Dürmuth. 2021. Apps Against the Spread: Privacy Implications and User Acceptance of COVID-19-Related Smartphone Apps on Three Continents. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, May 06, 2021. ACM, Yokohama Japan, 1–22. <https://doi.org/10.1145/3411764.3445517>
- [91] Jessica Vitak, Yuting Liao, Anouk Mols, Daniel Trotter, Michael Zimmer, Priya C. Kumar, and Jason Pridmore. 2023. When Do Data Collection and Use Become a Matter of Concern? A Cross-Cultural Comparison of U.S. and Dutch Privacy Attitudes. *International Journal of Communication* 17, 0 (January 2023), 28.
- [92] Jessica Vitak, Pamela Wisniewski, Xinru Page, Airi Lampinen, Eden Litt, Ralf de Wolf, Patrick Gage Kelley, and Manya Sleeper. 2015. The Future of Networked Privacy: Challenges and Opportunities. In *Companion of the 2015 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '15)*, 2015. ACM, New York, NY. Retrieved from https://cscw.acm.org/2015/program/accepted_workshops.php#W1
- [93] Jessica Vitak and Michael Zimmer. 2020. More Than Just Privacy: Using Contextual Integrity to Evaluate the Long-Term Risks from COVID-19 Surveillance Technologies. *Social Media + Society* 6, 3 (July 2020), 205630512094825. <https://doi.org/10.1177/2056305120948250>
- [94] Emily A. Vogels. 2020. About one-in-five Americans use a smart watch or fitness tracker. *Pew Research Center*. Retrieved January 14, 2023 from <https://www.pewresearch.org/fact-tank/2020/01/09/about-one-in-five-americans-use-a-smart-watch-or-fitness-tracker/>
- [95] Jacqueline Wernimont. 2018. *Numbered lives: life and death in quantum media*. The MIT Press, Cambridge, MA.
- [96] Pamela J. Wisniewski and Xinru Page. 2022. Privacy Theories and Frameworks. In *Modern Socio-Technical Perspectives on Privacy*, Bart P. Knijnenburg, Xinru Page, Pamela Wisniewski, Heather Richter Lipford, Nicholas

- Proferes and Jennifer Romano (eds.). Springer International Publishing, Cham, 15–41. https://doi.org/10.1007/978-3-030-82786-1_2
- [97] Philip Fei Wu, Jessica Vitak, and Michael T. Zimmer. 2020. A contextual approach to information privacy research. *Journal of the Association for Information Science and Technology* 71, 4 (April 2020), 485–490. <https://doi.org/10.1002/asi.24232>
- [98] Bin Xu, Pamara Chang, Christopher L. Welker, Natalya N. Bazarova, and Dan Cosley. 2016. Automatic Archiving versus Default Deletion: What Snapchat Tells Us About Ephemerality in Design. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, February 27, 2016. ACM, San Francisco California USA, 1662–1675. <https://doi.org/10.1145/2818048.2819948>
- [99] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata Mcdonough, and Yang Wang. 2019. Privacy Perceptions and Designs of Bystanders in Smart Homes. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (November 2019), 1–24. <https://doi.org/10.1145/3359161>
- [100] Shikun Zhang, Yan Shvartzshnaider, Yuanyuan Feng, Helen Nissenbaum, and Norman Sadeh. 2022. Stop the Spread: A Contextual Integrity Perspective on the Appropriateness of COVID-19 Vaccination Certificates. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, June 21, 2022. ACM, Seoul Republic of Korea, 1657–1670. <https://doi.org/10.1145/3531146.3533222>
- [101] Michael Zimmer. 2008. Privacy on Planet Google: Using the Theory of “Contextual Integrity” to Expose the Privacy Threads of Google’s Quest for the Perfect Search Engine. *Journal of Business & Technology Law* 3, 2 (2008), 109–126.
- [102] Michael Zimmer. 2018. Addressing conceptual gaps in big data research ethics: An application of contextual integrity. *Social Media + Society* 4, 2 (2018).
- [103] Michael Zimmer, Priya Kumar, Jessica Vitak, Yuting Liao, and Katie Chamberlain Kritikos. 2020. ‘There’s nothing really they can do with this information’: unpacking how users manage privacy boundaries for personal fitness information. *Information, Communication & Society* 23, 7 (2020), 1020–1037. <https://doi.org/10.1080/1369118X.2018.1543442>

Received January 2023; revised October 2023; accepted January 2024.