Contextual Integrity, Explained: A More Usable Privacy Definition

Nathan Malkin | University of Maryland and University of California, Berkeley

Contextual integrity is a privacy model popular with privacy researchers, but it can be relevant to others. This article introduces this theory and its main ideas, explains why it might be useful, and shows how it can be applied.

The theory of contextual integrity (CI)[1] provides a definition of privacy and a model for understanding when privacy violations happen. Since its introduction in 2004, it has become popular with privacy researchers, who have used it to explain why certain data practices have led to privacy controversies, predict when this might happen again, and understand people's privacy preferences. It can also be used by security and privacy workers hoping to answer similar questions. System architects and developers may wonder whether adding a new feature or using data in a novel way would create privacy problems. User experience designers and researchers may want to know which privacy choices require user attention and how to understand users' attitudes. Policy writers and implementers may want to figure out the best ways to protect citizens' privacy. More broadly, any one of us may wish to reflect on what it means to have privacy in the digital age, when living without generating vast amounts of data is not an option. CI can aid in answering all of these questions.

This article aims to serve as an accessible introduction to CI for researchers, practitioners, and others who have not encountered it before. It will argue for why a framework like CI is needed, describe the theory's major ideas, show examples of how it can be used, and discuss some of its limitations.

What Is Privacy? In Search of a Definition

Your privacy is very important. Everyone agrees about this. The United Nations Charter declares it to be a human right. Legislatures around the world pass laws to protect it. Companies announce their commitment to it in full-page ads (usually after violating it in some way). But what exactly do we mean when we talk about privacy?

Most people can readily come up with examples of behaviors they would consider privacy invasive: a peeping tom staring through a window, a stalker tracing a victim's whereabouts, an uninvited reader perusing a personal diary. But "I know it when I see it" is a cumbersome criterion by which to identify privacy violations. Moreover, cultures have different standards, and privacy preferences further differ between individuals.

When it comes to definitions, dictionaries are a natural place to seek clarity. Merriam-Webster, for example, defines *privacy* as "the quality or state of being apart from company or observation" or "freedom from

Digital Object Identifier 10.1109/MSEC.2022.3201585
Date of current version: 16 December 2022

unauthorized intrusion." While these definitions help clarify the notion of privacy, it's also apparent that they fail to cover a variety of situations and scenarios, especially when it comes to data and the digital domain. When we make a social media post, are we "apart from company and observation?" Why does some data usage feel creepy even when it is disclosed in terms of use documents? Can it still be considered an "unauthorized intrusion?" Dictionary definitions are too limited to provide insight into these questions and too vague to be operationally useful.

Legal definitions have the potential to be more specific, and laws such as the European General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) are specifically focused on regulating privacy in the Internet age. But while these laws define terms such as *personal data* and *aggregate consumer information*, they lack a succinct definition for the term *privacy*, instead codifying it as a series of rights for the data subject (such as the right to erasure and the right to rectification) and responsibilities for the data processor (to meet those rights). Moreover, just because some practice is legal doesn't mean people won't perceive it as a privacy violation, as numerous studies and media scandals attest.[2]

One of the main limitations of the definitions described so far is that they are difficult to use for analyzing situations. Concretely, researchers and practitioners in computer science often face questions about the privacy implications of a system:

- Does (or will) this system violate privacy?
- Does a solution preserve privacy or mitigate a privacy violation?

Ideally, a definition of privacy would provide enough insight to help address these questions. Essentially, we're looking for a model: something that can explain existing phenomena and be used to predict future outcomes. The theory of contextual integrity, invented and elaborated by Helen Nissenbaum,[1,3,4] offers just such a model for privacy. It can help analyze a situation from a privacy perspective and provide insights into how people will react when a new system or technology is introduced. CI has already been used successfully in a wide variety of computing research projects and beyond.[5,6,7,8] The rest of this article explains this theory and shows how to apply it.

CI: The Details

The theory of CI can be broken down into a few main ideas, each building on the previous ones. The theory is not all or nothing; you can adopt and use only some of the ideas while ignoring the rest.

Idea 1: Privacy Is Defined by How Information Flows

CI envisions privacy as the "appropriate flow" of personal information. The next section will define what it means for a flow to be appropriate, but first, let's take some time to explore why information flow is the most effective model for dealing with privacy.

Information flow refers to the transfer of knowledge from one party to the next. For example, when you report your symptoms to a nurse, who shares them with your doctor, who inputs them into a computer, which is then breached by a hacker, who sells the data on the illegal market—each of those points represents nodes through which your personal health information has flowed.

While the notion of information flow is fairly intuitive, CI emphasizes it because there are alternative models of privacy that are also widespread, for example, secrecy, data minimization, or leakage. The problem with these models is that they tend to be static and absolute: either something is secret, or it isn't. Some information might be classified as "sensitive" or "private"—for example, knowledge about relationships, finances, or health—leaving everything else to be labeled as "not sensitive," or maybe even "public."

CI, on the other hand, observes that privacy is fluid. As an illustration, consider that we don't hesitate to share gossip with our friends, financial records with our accountant, and health information with our doctor. But something would seem amiss if your friends started interrogating your tax returns, your accountant demanded a list of your medications, or your doctor insisted that you spill the latest gossip.

As this example shows, we can't divide information into "secret" and "not secret" or "private" and "public." Nor do friends, doctors, or accountants have "clearance" to access any of our sensitive details. In respective *contexts*, we freely share information we would otherwise consider private and off limits. Conversely, information that can be easily observed in public (such as a visit to a store and a purchase we make there) can still be considered private when it is taken out of the original context—for example, if it's aggregated to create a detailed profile of our movements or shopping habits.

CI addresses this problem by considering not only the specific data type but the information flow as a whole. Who were the intended recipients of the data, and what was their role? (See more about the details of the flow a bit later.) CI postulates that privacy violations happen when there is *inappropriate* information flow. But how do we distinguish appropriate and inappropriate information flows?

Idea 2: Information Flow Is Appropriate When It Conforms With Contextual Privacy Norms

According to the theory of CI, information flow is appropriate when it happens according to the norms of a particular informational context. In other words, CI asks, "What are the privacy norms in this specific situation?" If information is shared in a way that runs counter to these entrenched expectations, that flow is inappropriate—i.e., it is a privacy violation. In fact, this is precisely how the theory defines privacy:

> Privacy, defined as CI, is preserved when information flows generated by an action or practice conform to legitimate contextual informational norms; it is violated when they are breached (Nissenbaum,[4] p. 224).

While this may appear almost tautological ("a privacy violation happens when you violate privacy expectations"), this definition draws an important distinction from notions of privacy that are purely procedural, such as the principle of informed consent and other Fair Information Practice Principles (FIPPs). Under a procedural model of privacy, any information flow might be considered appropriate as long as certain practices were followed, such as encrypting the data in transit or getting the user to agree to some terms and conditions.

Informed consent and other FIPPs certainly have their value, but CI says that following them is not sufficient to maintain privacy, just like you're unlikely to achieve security simply by ticking all the boxes on a checklist. Privacy concerns won't magically go away just because the user clicked "I accept." Instead, CI postulates that norms are the key determinant for privacy.

Norms are generally established standards and commonly held expectations about what will happen with shared information. Here are some examples of contextual informational norms:

- A teacher is expected to share a pupil's grades with the student's guardians (and perhaps other teachers) but not anyone else.
- A therapist is expected not to reveal a patient's mental state unless they believe the patient is in danger.
- Citizens are required to report their income to the government, but the government is expected not to make that information public.

As these examples illustrate, norms are like the rules that govern our interactions in society. Some may be informal but enduring (for example, norms about sharing intimate details or betraying a friend's confidence). They can be so strongly held that they have been codified as laws (such as, in the United States, privacy regulations about health data or children). But they might also be loosely defined and best-effort (getting our friends' permission before photographing them or posting those pictures on social media). Finally, some may be vague and rapidly evolving (such as the question of how to respect the privacy of guests in smart homes).

Just like social rules, norms can be shared by an entire society or country (for example, being obligated to submit one's fingerprints if arrested on suspicion of a crime) or can be localized to an individual family or workplace (such as a company where all employees know each other's compensation). To summarize:

> Norms may be explicit or implicit, 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 (Nissenbaum,[4] p. 227).

An implication of this flexibility is that there may not be a single true norm for a given situation. Multiple norms may be present, and perhaps even in conflict with each other, due to the interaction of different contexts, cultures, and values. As a result, not everyone might agree about what the prevailing norm is. In these cases, CI doesn't necessarily offer a resolution (though it provides some guidance, to be discussed in a later section), but it can help model what is happening.

Nonetheless, the flexibility of norms is not total. Since they are like social rules and represent entrenched expectations in society, norms require some consensus: one person's opinion, no matter how reasonable or well-justified, cannot constitute a norm if it is at odds with everyone else's.

Therefore, norms cannot be assumed or derived from first principles but must rather be gleaned from the real world. Thus, the most reliable way to ascertain a norm is to identify people's attitudes, beliefs, and expectations. Because norms differ between contexts, conducting this research (and, more generally, understanding norms) requires a more precise definition of what constitutes a context.

Idea 3: A Contextual Norm Can Be Described by (at Least) Five Parameters

So far, we've seen that privacy can be modeled as information flow and argued that the privacy expectations for these flows are governed by norms, which vary according to context. But what exactly constitutes a context?

According to the theory of CI, a context can be defined by the following parameters:

1. data type (what sort of information is being shared)
2. data subject (who the information is about)
3. sender (who is sharing the data)
4. recipient (who is getting the data)
5. transmission principle (the constraints imposed on the flow).

To figure out the privacy norms at play in a particular situation, you need to identify and consider all five of these variables (Figure 1).

According to CI, if one of these variables is undefined, the situation is underspecified, and the privacy expectations can't be fully determined. For example, if we don't know what the information is or whom it's about, we can't say how it should be shared. Or if we know those things but we don't know whom it's being shared with, we don't know if privacy violations are occurring.

Data type, subject, sender, and recipient are all fairly self-explanatory; they've already been implicit in our discussion of information flow. The transmission principle parameter is new to CI and therefore requires some elaboration.

The *transmission principle* accounts for the conditions or constraints that restrict information flow or limit it to specific circumstances. For example, according to some norms, a business should share its customers' records with the government only if the authorities have a warrant or court order. Here, the transmission principle is the existence of a warrant; only in its presence does the information flow become appropriate.

Other potential transmission principles include:

- the subject's consent
- the consent of a parent or guardian (usually when the subject is a minor)
- with notice (some sort of advance announcement or disclosure)
- reciprocity ("I'll show you mine if you show me yours")
- subject to legal requirements
- the Chatham House Rule (information can be reshared only without attribution).

This list is far from exhaustive; there are many other transmission principles.

There may also be other CI parameters. While CI holds that the five variables (data type, subject, sender, recipient, and transmission principle) are generally sufficient for specifying a context, it allows that other factors may influence people's expectations and norms.

One specific example that often comes up is the question of the purpose or use (that is, how some data will be used and to what end). This turns out to be an important factor both from a legal point of view and in people's expectations.[9] For example, smart speaker users share their voice and interaction data with voice assistants, expecting that these will be used to answer queries, provide services, and perhaps improve the devices; however, many would find it unacceptable if these data were used for advertising.[8] This distinction could be represented by a separate "purpose" parameter.

The CI model, in its original formulation, lacks this purpose/use variable, though Nissenbaum, the theory's creator, has written that she is "increasingly persuaded" that it should be included[4] (p. 234). However, CI does provide a framework for addressing this distinction. CI conceives of actors (subjects, senders, and recipients) not as identities (named individuals and companies) but as *roles* (capacities in which they act). An actor might have different roles; for example, your doctor might happen to be your friend or family member. In that case, privacy norms are determined by that person's role in a particular context: if they receive information in their capacity as a health-care provider, expectations are different than if they had heard the same thing at a family function.

Roles can be used to specify and restrict purpose. Returning to the question of smart speaker users, we can say that they are sharing their data with voice assistant companies in their role as information providers. If those companies use it for advertising, then they are taking on a different role—that of advertisers—which is outside the expected context.

Regardless of how exactly you choose to model context, it's worth remembering that purpose matters and that there can be more to CI than just the five parameters.

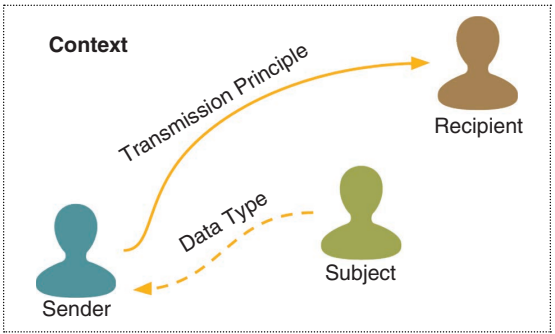

Figure 1. The five parameters defining context according to the theory of CI. Note that Nissenbaum[4] (p. 227) emphasizes that "respective roles, activities, purposes, information types do not exist *in a* context; rather, these factors *constitute* a context."

Idea 4: New Norms and Flows Are Evaluated Through Their Context

The previous ideas have described CI's perception of privacy and conception of contexts; together, these provide us with the ability to model existing information flows with respect to the privacy norms that govern them. But what happens if there's a new information flow?

Just because an information flow is unfamiliar, doesn't mean it's bad; the new flow could still be appropriate. However, when we're dealing with new technology and novel data flows, norms for a specific context might not be established yet.

Consider, for example, a doctor who wants to make use of new technologies that require access to patients' data (perhaps AI-powered dictation software, a diagnostics assistant, or an electronic health record). These are novel flows, so how can the doctor determine whether or not they are appropriate?

The theory of CI provides a way to evaluate the *ethical legitimacy* of new flows. It gives a framework for identifying the strengths and weaknesses of the novel flow as compared with the status quo. CI suggests three layers of analysis:

1. the interests of the affected parties
2. the ethical and political values
3. the contextual functions, purposes, and values.

In the case of the doctor's data sharing, we would first consider the interests of the parties. It will make the doctor's life easier, but are there ways in which it might be detrimental to the patients? Next, we would look at more general ethical priorities—for example, values like justice and equity, free speech, and freedom of choice. Would any of these be hurt? Finally, we'd think about the fundamental purpose of the context—in this case, providing health care. Would its goals be undermined by the flow, or any of its consequences? For example, will patients become less likely to seek care due to concerns about how their data are used?

Each of these factors may offer a reason to reject a new flow as not being morally legitimate from the perspective of CI. If the data sharing leads to higher insurance premiums, it can hurt the patient financially. If it results in disparate health outcomes for different demographics, it may be unjust. On the other hand, the benefits offered by the new flow might show it to be superior to the status quo. If the data sharing aids the contextual purpose of providing health care—for example, by stopping an outbreak of a disease early through notifying public health officials so they can take appropriate measures—this, according to CI, could outweigh the privacy interests of an individual data subject.

Clearly, these determinations are subjective and might be contested; even outside the realm of privacy, debates rage over whether particular policies align with specific values. While CI cannot deliver a definitive decision in each case, the framework provides a structured way of thinking about whether something hurts or enhances privacy.

Applying CI

This article has argued that CI offers an effective model and a more precise definition of privacy. But how can it be used? Next are four lessons demonstrating how CI can be applied to research and practice.

Lesson 1: Think Beyond Binaries

It can be tempting to reduce data to binary categories: sensitive or not sensitive; private or public, information that is—or isn't—personally identifiable, and so on. Yet, just as anonymous data can often be reidentified, so can public data often turn out to be sensitive. All of these binary characterizations fail to acknowledge the context-dependent nature of what people consider private.

Of course, this isn't a suggestion to start treating credit card numbers the same way as comments on a blog post. If anything, it's the opposite. For example, if one were to aggregate a person's every public comment and product review into a dossier and then publish it, that would feel like a privacy violation. Why? Weren't they public already? As CI explains, it's not enough to consider that the information is public; we need to think about how that information was flowing before and how that flow changed.

Another illuminating example is the outcry when, in short succession, pretty much every voice assistant was revealed to have been relying on contractors to listen to some user interactions.[10] Many people were upset to discover this new, previously undisclosed, flow, forcing the companies to apologize and backtrack.

In this situation, the companies felt that they were relatively unconstrained by what they could do with the data since users had already shared those recordings with them. In reality, they were taking interactions that many saw as ephemeral and generating new data flows on their basis, creating a (mostly invisible) permanent record. The companies consequently learned that the new flows were surprising and unwelcome to people even though the data technically never left the company and was not shared with third parties. These scandals may have been avoided had the companies been thinking in terms of information flows and also if they had checked how any such new flows aligned with people's expectations.

Lesson 2: Check Expectations, Not Checklists

Internet history is replete with services that abused their users' trust and data and then pointed to a line of fine print to justify it: "Can't you see? You agreed to all of this." Courts have been increasingly skeptical of this defense, and CI explains why it was never satisfactory. What we consider to be a privacy violation is based on our expectations for a particular context, not a set of practices the provider did or didn't follow.

Newer legal frameworks, such as GDPR and CCPA, are recognizing this and are consequently requiring positive assent with meaningful opt-out options instead of pro forma checkboxes that everyone has to click through. Other pro-privacy moves can also be necessary but not sufficient. For example, data minimization, while a positive step, may not, on its own, be enough to assuage privacy concerns.

Even privacy-enhancing technologies (PETs) can fall short due to a mismatch in consumer expectations. For example, research found that many users misunderstood web browsers' private browsing modes, thinking that their browsing history would be secret from entities such as employers, governments, or Internet service providers.[11]

As discussed previously, this is not a dismissal of practices like data minimization or informed consent. They are useful tools on the path to privacy—the path, that is, to following people's expectations and adhering to norms.

What are those expectations? The easiest way to find out is to ask. Researchers in a number of academic fields (anthropology, sociology, information science, and human–computer interaction) have been studying these questions for years and have developed techniques for discovering user expectations in general and for CI specifically.[6] Similar methods are also used daily by user experience researchers in industry, who are working in large numbers at companies big and small.

Lesson 3: Account for the Complete Context

One important thing to remember about expectations is that they are specific to contexts. Therefore, just because something is considered acceptable in one context doesn't mean it'll be okay in another. For example, social media buttons ("Like this! Share that!") are considered acceptable on news and lifestyle websites but raise questions when they appear on health sites. Though the data flows are ostensibly similar, the contrasting contexts mean the expectations are different.

To think through a context and consider ways in which it might differ from more familiar ones, it can help to identify the parameters singled out by CI: data type, subject, sender, recipient, and transmission principle. If even just one of these variables changes—for example, a new recipient is added or a transmission principle such as reciprocity is lacking—then the entire flow may become inappropriate.

The details of the parameters matter. Returning to the example of human review of voice assistant recordings, we can reason that users may have known their recordings were being sent to the company. However, they likely assumed that their recordings were being processed algorithmically and were never exposed to other people. Established norms did not account for the listening by human beings, even if it was done for benign purposes like improving the assistants' performance. In general, research has found that people are wary of their data being examined by humans (as opposed to being processed automatically by machines) and of that data being shared with third parties, whether for advertising or other purposes.[7]

The details of information flows are relevant to privacy-enhancing technologies as well because they may inadvertently introduce new flows. For example, when web browsers introduced the Do Not Track HTTP header, it was intended for users to signal an opt-out from behavioral advertising, but it actually ended up being used as another signal for fingerprinting browsers and tracking users.[12]

Examples like these provide an important reminder that, when introducing changes to a sociotechnical system, we need to verify the contextual integrity of the proposed system:

- Will new information flows be introduced?
- Are existing information flows changing?
- What are the effects of these changes?

The latter question—the consequences of privacy changes—is especially crucial to consider.

Lesson 4: Consider the Consequences

As we have seen, CI can help understand the privacy implications of new technologies by decomposing novel information flows into their constituent components (data type, subject, and so on). However, the CI framework is also helpful for higher-level reasoning about privacy. This is enabled by the theory's focus on contextual purposes.

Why do we share information with other people? Usually, the information flow serves a specific goal. Data are shared in medical contexts for the purpose of curing patients, in education contexts for the purpose of imparting knowledge to students, and in contexts of the judicial system for the purpose of securing justice. Even casual interactions, like gossip or small talk, serve some social motive.

CI instructs us to consider the consequences for these purposes when analyzing the impact of new flows. This framework can be used, as an example, to analyze the concerns surrounding the surveillance of students. As part of the pandemic-induced switch to remote learning, students were subjected to a variety of new demands on their privacy, from requirements to turn on their webcams and be on video during remote lectures to invasive monitoring of their computers and surroundings as a part of remote proctoring.[13] How should we think about the ethical legitimacy of these novel flows?

The rights and interests of students and instructors are a good starting point for this debate. But CI offers an additional question to guide deliberation. Do any of these measures enhance student learning? Or do they actually hurt students' education by drawing their attention away from the subject matter and introducing new stresses? If so, then the new flows are privacy violations and inappropriate.

Similar skepticism should be shown to new flows that endanger the values and purposes of other contexts: health technologies that may make patients reluctant to seek care (for example, data sharing between health-care providers and employers) and voting methods that may reduce citizens' engagement or increase their distrust in civic affairs (such as certain proposals for online voting). Regardless of the setting or the technology, a full appraisal calls for considering the contextual values.

Ultimately, this perspective is so useful because—just as security is not a primary activity but rather an operational requirement—most people don't care about privacy for its own sake. Privacy enables free speech, creativity, self-expression, experimentation, and other beneficial values and outcomes. When we fight for privacy, we fight for these values, too.

Unresolved Questions and Future Directions

Beyond the lessons discussed previously, there may be opportunities to incorporate CI more directly into the privacy decision-making of systems. Exactly how this might be done remains an ongoing research question. In the meantime, the theory has already proven useful in a number of computer science research projects.[5] Still, CI isn't the final word in privacy; it has a number of limitations, which are worth knowing about.

As a theoretical model, CI aims to predict how people will feel about privacy under different circumstances; it does not claim that this is how people think about privacy or make privacy decisions. You're unlikely to find many people who go into a situation, identify each of the five CI parameters at play, reflect on the context they are operating in, and then arrive at a privacy judgment. Most of the time, our reactions are rooted in emotions and intuitions.

Furthermore, even if asked to reflect more logically on their decisions, people don't necessarily think about the situation in the same terms as the CI model.[14] And like any model, CI necessarily simplifies things. As discussed previously, there may be other factors that matter, beyond the parameters CI identifies.

Another limitation of CI is its conservativeness. Though it provides a way of adjudicating novel flows based on the moral values at play, CI favors established norms. Existing expectations can be entrenched for good reasons, but not always. For example, many workplaces have a norm that employees don't share their salaries with each other, but this may have the effect of limiting workers' bargaining power and hurting underrepresented minorities. CI provides some tools for reasoning about these disputes, but it's not a complete theory of norm evaluation.

One of the biggest challenges for CI is the problem of inferences, in which the collection of one data type can lead to conclusions about another, as a result of which harmless data bits can be composed into highly invasive profiles. A famous example is a woman's pregnancy that was predicted, based on shopping history, prior to her knowing.[15] CI's perspective on this is that higher-order data types—ones derived and inferred from other information—should be evaluated on their own terms, not based on the norms of the lower-order source data. Just because it's accepted for a store to keep track of your purchases doesn't make it okay for it to traffic in health data that they were able to infer from your buying habits. If anything, privacy expectations might "travel down": If some data can be used to infer something sensitive, then that original information should be subject to the same constraints as the inferred sensitive data would be.

The challenges of inferences are threefold. Data primitives—such as electric impulses, clicks, and page views—lack meaningful privacy semantics on their own. It can be difficult to predict how they (and other lower order data types) will compose to become more complex information with privacy implications. The rapid pace of technological change means that new techniques and possibilities for inferences emerge regularly. Because norms may take time to become established, the result is that privacy rules can struggle to keep up. These are open questions not only for CI but for other conceptualizations of privacy. The problem isn't just theoretical; in a world of big data, inferences can pose as much of a privacy threat as direct observations. Solutions—both for theory and in practice—are urgently needed.

Contextual integrity, like our understanding of privacy more generally, continues to evolve, and in time, a new model or an improved definition might come along to extend (or even replace) this theory. But CI is already a powerful tool for making sense of and helping ensure privacy. As researchers and practitioners in computer science, everyone would benefit if more of us knew about and made use of CI. ■

Acknowledgment

Thanks to Florian Schaub, David Wagner, and Helen Nissenbaum for their comments on draft versions of this article. Any remaining errors are the author's own.

References

1. H. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA, USA: Stanford Law Books, 2009.
2. K. Sarikakis and L. Winter, "Social media users' legal consciousness about privacy," *Social Media + Soc.*, vol. 3, no. 1, p. 2,056,305,117,695,325, 2017, doi: 10.1177/2056305117695325.
3. H. Nissenbaum, "Privacy as contextual integrity," *Washington Law Rev.*, vol. 79, p. 119, Feb. 2004.
4. H. Nissenbaum, "Contextual integrity up and down the data food chain," *Theor. Inquiries Law*, vol. 20, no. 1, pp. 221–256, 2019, doi: 10.1515/til-2019-0008.
5. S. Benthall, S. Gürses, and H. Nissenbaum, "Contextual integrity through the lens of computer science," *Found. Trends Privacy Secur.*, vol. 2, no. 1, pp. 1–69, 2017, doi: 10.1561/3300000016.
6. Y. Shvartzshnaider et al., "Learning privacy expectations by crowdsourcing contextual informational norms," in *Proc. 4th AAAI Conf. Human Comput. Crowdsourcing*, 2016, pp. 1–10.
7. N. Apthorpe, Y. Shvartzshnaider, A. Mathur, D. Reisman, and N. Feamster, "Discovering smart home Internet of Things privacy norms using contextual integrity," *Proc. ACM Interact. Mobile Wearable Ubiquitous Technol.*, vol. 2, no. 2, pp. 59:1–59:23, Jul. 2018, doi: 10.1145/3214262.
8. N. Abdi, X. Zhan, K. M. Ramokapane, and J. Such, "Privacy norms for smart home personal assistants," in *Proc. 2021 CHI Conf. Hum. Factors Comput. Syst.*, New York, NY, USA: Association for Computing Machinery, pp. 1–14, doi: 10.1145/3411764.3445122.
9. P. Emami-Naeini et al., "Privacy expectations and preferences in an IoT world," in *Proc. 13th Symp. Usable Privacy Secur. (SOUPS)*, Santa Clara, CA, USA, 2017, pp. 399–412. [Online]. Available: https://www.usenix.org/conference/soups2017/technical-sessions/presentation/naeini
10. M. Day, G. Turner, and N. Drozdiak. "Amazon workers are listening to what you tell Alexa." Bloomberg. Accessed: Nov. 1, 2022. [Online]. Available: https://www.bloomberg.com/news/articles/2019-04-10/is-anyone-listening-to-you-on-alexa-a-global-team-reviews-audio
11. X. Gao, Y. Yang, H. Fu, J. Lindqvist, and Y. Wang, "Private browsing: An inquiry on usability and privacy protection," in *Proc. 13th Workshop Privacy Electron. Soc.*, New York, NY, USA, 2014, pp. 97–106, doi: 10.1145/2665943.2665953.
12. N. Nikiforakis, A. Kapravelos, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna, "Cookieless monster: Exploring the ecosystem of web-based device fingerprinting," in *Proc. IEEE Symp. Secur. Privacy*, May 2013, pp. 541–555, doi: 10.1109/SP.2013.43.
13. J. Kelley. "Students are pushing back against proctoring surveillance apps." EFF Deeplinks Blog. Accessed: Nov. 1, 2022. [Online]. Available: https://www.eff.org/deeplinks/2020/09/students-are-pushing-back-against-proctoring-surveillance-apps
14. A. Frik, J. Bernd, N. Alomar, and S. Egelman, "A qualitative model of older adults' contextual decision-making about information sharing," in *Proc. Workshop Econ. Inf. Secur. (WEIS)*, 2020, pp. 1–62.
15. C. Duhigg. "How companies learn your secrets." NY Times. Accessed: Nov. 1, 2022. [Online]. Available: https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html

NATHAN MALKIN is currently a postdoctoral researcher at the University of Maryland, College Park, MD 20742 USA. His research interests include usable security and privacy, most recently developing and testing novel approaches to managing privacy in smart homes, and he has found contextual integrity invaluable in his work. Malkin received his Ph.D. in computer science from the University of California, Berkeley. Contact him at nmalkin@cs.berkeley.edu.