Privacy Mini-Publics: A Deliberative Democratic Approach to Understanding Informational Norms

Daniel Susser

Department of Information Science, Cornell University

Matteo Bonotti

Department of Politics and International Relations, Monash University

Nissenbaum's theory of contextual integrity (CI) provides a powerful lens through which to conceptualize, discuss, and advocate for information privacy, and it has quickly become a central tool for structuring theory and practice in philosophical, legal, policy, and engineering domains. However, as an abstract analytical model, CI does not offer substantive details about the actual prevailing privacy norms operative in any given context. An impressive body of empirical research has emerged to fill this gap, using interview and survey methods to discover context-relative privacy expectations. Yet we argue that such approaches are limited by a fundamental mismatch between method and object of study. According to CI, privacy norms are social, *collective* norms—survey methods attempt to infer such norms, indirectly, by measuring, aggregating, and analyzing individual expectations and preferences. While such approaches have produced promising insights, we propose a methodology for surfacing privacy norms—privacy mini-publics—that centers their collective nature and leverages a collective, deliberative process to understand them. We expect this methodology to yield both epistemic and political advantages over the state of the art, generating more accurate and widely accepted results. If our experiments with privacy mini-publics prove fruitful, we suggest integrating them with survey-based methods to produce a more robust, hybrid approach.

CCS Concepts • Insert your first CCS term here • Insert your second CCS term here • Insert your third CCS term here

Additional Keywords and Phrases: Privacy, Contextual Integrity, Deliberative Democracy, Mini-Publics, Methodology

ACM Reference

Daniel Susser and Matteo Bonotti. 2023. Privacy Mini-Publics: A Deliberative Democratic Approach to Understanding Informational Norms. In PrivaCI '23: Symposium on Applications of Contextual Integrity, September 23-24, 2023, Toronto, Canada.

1 INTRODUCTION

Nissenbaum's theory of contextual integrity (CI) provides a powerful lens through which to conceptualize, discuss, and advocate for information privacy. In contrast with earlier approaches, which treated privacy as everywhere and always one and the same thing—privacy as secrecy (Posner 1978), for example, or privacy as "control over personal information" (Westin 1967)—contextual integrity understands information privacy as a set of diverse, context-specific social norms (Nissenbaum 2010). In one setting, say the doctor's office, expectations of privacy might demand giving data subjects control over information about themselves, as control theories suggest. But in other settings, such as the home or the classroom, expectations may differ, allowing for certain information to be shared with certain other recipients (between parents, perhaps, or between teachers and student advisors). CI captures these nuances by modeling "context-relative informational norms" (i.e., privacy norms) in terms of five parameters: data subjects, senders, recipients, types, and transmission principles. And it enables rigorous analysis of the disruptions to prevailing privacy norms that digital technologies threaten, by providing a framework for carefully describing the changes such technologies make to the ways information flows.

Given the precision CI brings to discussions about information privacy, it has quickly become a central tool for structuring theory and practice in philosophical, legal, policy, and engineering domains. Yet these endeavors inevitably face one important obstacle, which CI (as an abstract analytical model) cannot, by itself, overcome—namely, identifying the actual, prevailing context-relative informational norms operative in a given context (i.e., the starting point or baseline for evaluating information flows). If policymakers are going to create legal protections for information privacy, and if engineers are going to develop systems that respect it, they need to know not just how to think about and conceptualize privacy; they need to know what our actual privacy norms *are.* CI provides a framework for investigating such norms, but it does not provide substantive answers to the question: "What are the social expectations for how information should flow in this context?"

Answering that question requires empirical research, and a growing body of scholarship attempts to undertake it—exploring privacy norms in the contexts of, amongst other things, healthcare (e.g., Nicholas et al. 2019), education (e.g., Jones et al. 2020), and the home (e.g., Apthorpe et al. 2018). Methodologically, these studies predominantly rely on approaches from user experience (UX) and human-computer interaction (HCI) research, such as interviews and surveys (Badillo-Urquiola et al. 2019). It is common for researchers to poll people's intuitions via Amazon Mechanical Turk (AMT) or similar crowd-sourcing systems, using sophisticated factorial vignette surveys that vary prompts in order to gauge subtle differences in privacy expectations based on the specific senders, recipients, data types and other features of information flows in particular contexts (e.g., Martin and Nissenbaum 2015; Shvartzshnaider et al. 2016).

This work has yielded important insights into real-world privacy norms, giving researchers and practitioners guidance about how to design privacy protective systems and helping advocates establish normative baselines from which to diagnose worrying technological disruptions. However, these approaches suffer from an important limitation: fundamentally, they attempt to infer social norms from the opinions, preferences, and attitudes of individuals. Which is to say, they try to answer the question "What are the social expectations for how information should flow in context X?" by posing the question "How should information flow in context X?" to a variety of people (perhaps even a random, representative sample of some population), in isolation, and analyzing the aggregate results. That is not an unreasonable strategy, since presumably there is some relationship between individual preferences and social norms, and researchers have developed ingenious methods for inferring social expectations from individual ones. Shvartzshnaider et al. (2016), for example, attempt to discover an "implicit consensus" about social norms by measuring both the preferences of individual survey respondents and the degree to which individual responses diverge from those held by the majority.

This project proposes a different methodology for investigating contextual integrity norms, one that *starts* from the collective nature of social norms, rather than trying to reach it indirectly via aggregates of individuals. Drawing from empirical research in political science, we propose the use of "deliberative mini-publics"—forums of randomly selected citizens who deliberate, collectively, on a policy issue and arrive at recommendations aimed at influencing public policy (Smith and Setälä 2018)—to identify and understand informational norms in particular contexts. That is, instead of asking individuals to reflect on their own expectations or preferences, privacy mini-publics would put groups of people into structured conversations and ask them to reach (explicit) consensus about social expectations of privacy. We anticipate that such an approach will have both epistemic and political advantages over the state of the art, more accurately characterizing prevailing norms and guiding the development of more democratically legitimate information systems.

2 THE METHOD

This project adopts an innovative approach centred around deliberative mini-publics, drawing from best-practice principles for deliberative processes, as described by the Organisation for Economic Cooperation and Development (OECD) in the report 'Innovative Citizen Participation and New Democratic Institutions—Catching the Deliberative Wave' (OECD 2020). Of particular importance are the following OECD principles, which inform each stage of our mini-pubic design:

2.1 Stage 1 (Recruitment)

> 'Participants should be a microcosm of the general public' (OECD 2020, p. 16).

Each mini-public will consist of a forum of approximately 40 US citizens, selected through stratified random sampling (sortition) to broadly reflect local demographics (e.g. age, gender, location, cultural background, employment types and household environment). This recruitment process is crucial for deliberation because it helps to guarantee a diverse and inclusive group in which a variety of views are represented. The recruitment of each mini-public will be conducted by an organization such as the Sortition Foundation.

2.2 Stage 2 (Experts and Information)

> 'Participants should have access to a wide range of accurate, relevant and accessible evidence and expertise' (OECD 2022, p. 17).

To ensure that participants are meaningfully informed about relevant technical details, policy frameworks, and interests involved in particular information flows, they will be briefed by experts of three kinds:

1. Academic experts—such as philosophers, ethicists, political scientists, and lawyers—whose knowledge comes from study, research and professional practice in the area of privacy;
2. Government representatives—such as legislators and civil servants—whose knowledge comes from lived experience related to policy-making in the area of privacy;
3. Industry representatives—such as privacy/data protection officers or ethics and compliance professionals—who are charged with helping technology firms navigate privacy laws, norms, and expectations.

2.3 Stage 3 (Deliberation and Facilitation)

> 'Group deliberation entails finding common ground; this requires careful and active listening, weighing and considering multiple perspectives…and skilled facilitation' (OECD 2022, p. 17).

Following the expert briefing, the participants will engage in a deliberative exchange for up to 20 hours over four sessions, during which they will consider the contrasting viewpoints on the issue at stake, identify the preferred outcome for their community, and work through the trade-offs that arise during this discussion. The goal of the deliberative exchange will be to reach consensus on (1) what the existing social norms about privacy are in a specific context; (2) whether the introduction of some new data-driven technology in that context has disrupted such norms; and (3) what kinds of interventions and regulations (if any) are required to address these changes. As recent work in the field has pointed out, consensus, strictly defined, 'has, for the most part, been abandoned by deliberative democrats' (Dryzek 2016, p. 230). Therefore, for the purposes of this project and, again, reflecting its practical orientation, we understand 'consensus' as an agreement at the level of 80% or more.

During the deliberative sessions, participants will be encouraged to actively listen, empathize with other standpoints and engage with a wide range of views about the policy issue. Expert facilitators from democracyCo[1] (or similar organization) will guide each session and encourage participants to apply critical thinking practices to reduce unconscious biases. The sessions will include small-group work to enable in-depth and inclusive deliberation, with a focus on discussion questions that facilitators prepare in order to open up debates and enable comparison between groups during the plenary discussions that follow. Participants will also be encouraged to use the time between sessions to further reflect on what they have learned. The recommendations agreed to by the participants at the end of the deliberative process will be compiled in a report.

3 DISCUSSION

The aim of this proposal is to introduce a new method for empirically discovering context-relative informational norms, which complements existing approaches. Recognizing and foregrounding the collective nature of social norms, privacy mini-publics would endeavor to understand the specific norms operative in a given context by directly leveraging a collective, deliberative process, rather than attempting to infer collective norms indirectly via prompts to individuals.

[1] https://www.democracyco.com.au/

Deliberative democratic theory promises both epistemic and political improvements over aggregative approaches. Deliberation is thought to yield *better* (i.e., truer or more just) and more *legitimate* (i.e., widely accepted) political insights and decisions, for a number of reasons: In contrast with individual reflection or introspection, intersubjective modes of reasoning are likely to introduce more information and evidence and to offer a greater range of perspectives through which to interpret and understand it. Cognitive diversity—i.e., differences in "the way people think about a problem in the world"—generates more creative and effective problem solving. Disagreement invites participants to critically evaluate their beliefs and preconceptions, and the demand for consensus encourages participants to revise them (Estlund and Landemore 2018, p. 120-3).

We expect privacy mini-publics to offer similar advantages over the state of the art in empirical privacy scholarship. Epistemically, we anticipate that the deliberative mini-public process will surface privacy norms more accurately than survey-based methods. Further research will be required to evaluate this hypothesis—we propose to construct the first privacy mini-publics around scenarios familiar from existing studies, and then to use survey methods to compare their results. Politically, we anticipate the process of engaging in deliberation to produce more legitimate results—i.e., more consensus among participants about not only what prevailing privacy norms are and whether new technologies threaten to disrupt them, but also about how to mitigate or manage these changes.

At the same time, mini-publics suffer from an important limitation: they do not scale. Designing a mini-public, recruiting experts and participants, carrying out many hours of deliberation, and compiling results is a complex, time-consuming, and costly endeavor. By contrast, survey-based methods can be scaled more easily, and efforts are underway to develop methods for automating the evaluation of their results. Given the vast number of privacy-related decisions that need to be made, every day, in industry, government, and elsewhere, flexible and scalable approaches are necessary. Thus our goal is not to replace current methods, but rather to build upon them: if our experiment with privacy mini-publics proves fruitful, we hope to integrate them into a more robust, hybrid approach.

REFERENCES

Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering Smart Home Internet of Things Privacy Norms Using Contextual Integrity. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 2, 2, Article 59.

Karla Badillo-Urquiola, Xinru Page, and Pamela Wisniewski. 2018. Literature Review: Examining Contextual Integrity within Human-Computer Interaction. *Available at SSRN 3309331*.

John Dryzek. 2016. Symposium commentary: Reflections on the theory of deliberative systems. *Critical Policy Studies*, 10, 2, 209-215.

David Estlund and Hélène Landemore. 2018. The Epistemic Value of Democratic Deliberation. In *The Oxford Handbook of Deliberative Democracy*. Oxford University Press.

Kyle ML Jones, Andrew Asher, Abigail Goben, Michael R Perry, Dorothea Salo, Kristin A Briney, and M Brooke Robertshaw. 2020. "We're being tracked at all times": Student perspectives of their privacy in relation to learning analytics in higher education. *Journal of the Association for Information Science and Technology*.

Kirsten Martin and Helen Nissenbaum. 2015. Measuring privacy: an empirical test using context to expose confounding variables. *Columbia Science and Technology Law Review* 18, 176.

Jennifer Nicholas, Katie Shilton, Stephen M Schueller, Elizabeth L Gray, Mary J Kwasny, and David C Mohr. 2019. The Role of Data Type and Recipient in Individuals' Perspectives on Sharing Passively Collected Smartphone Data for Mental Health: Cross-Sectional Questionnaire Study. *JMIR mHealth and uHealth* 7, 4.

Helen Nissenbaum. 2010. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.

OECD [Organisation for Economic Co-operation and Development]. 2020. *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave*. https://www.oecd.org/gov/open-government/innovative-citizen-participation-new-democratic-institutions-catching-the-deliberative-wave-highlights.pdf

Richard Posner. 1978. The Right of Privacy. *Georgia Law Review* 12, 3, 393.

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*.

Graham Smith and Maija Setälä. 2018. 'Mini-Publics and Deliberative Democracy', in A. Bächtiger, J.S. Dryzek, J. Mansbridge and M. Warren (eds.), *The Oxford Handbook of Deliberative Democracy*. Oxford University Press, pp. 300-14.

Alan Westin. 1967. *Privacy and Freedom*. Atheneum, New York.