

Introduction

Madelyn Rose Sanfilippo, Brett M. Frischmann, and Katherine J. Strandburg

Privacy, in contrast with secrecy, is a relational concept, achieved when personal information is shared appropriately between actors. Viewed in this way, privacy is necessarily contextual and complex because norms about appropriate flows and use of personal information are socially negotiated and often contested (Nissenbaum, 2009). Privacy is thus a problem of collective action. Moreover, personal information is often among the knowledge resources pooled and managed by knowledge commons. Even when that is not the case, personal information can be important in shaping knowledge commons participation and governance. The Governing Knowledge Commons (GKC) framework is thus well suited for studying and analyzing how communities or populations evaluate and shape governance of privacy in particular contexts. (Sanfilippo, Frischmann & Strandburg, 2018)

Chapter 1 of this volume introduces the theoretical basis for applying the GKC framework to study privacy, explores how that framework complements and supplements Nissenbaum's contextual integrity theory, and describes a privacy-focused meta-analysis of previous GKC case studies. Previous case studies within the GKC tradition did not explicitly address questions of privacy. Nonetheless, the meta-analysis presented in Chapter 1 demonstrates that personal information shaped governance – and was itself pooled and governed – in previously published GKC cases. By studying how the strength and enforcement of particular types of “rules-in-use” for personal information varied among those cases, the privacy-focused meta-analysis uncovers three patterns of commons governance: member-driven, public-driven, and imposed.

Drawing on insights from the theory and meta-analysis reviewed in Chapter 1, the chapters gathered in this volume were solicited from an interdisciplinary group of scholars studying personal information governance in a variety of contexts. Chapters 2 through 5 in this volume present case studies of knowledge commons in which personal information is pooled and governed as a critical knowledge resource. Chapters 6 and 7 present case studies in which privacy's role is primarily instrumental to the creation and management of other sorts of knowledge resources; commons

2 *Madelyn Rose Sanfilippo, Brett M. Frischmann, and Katherine J. Strandburg*

governance of personal information enables trust and cooperation. Chapters 8 through 10 explore some of the failures and complexities of privacy commons governance, particularly with respect to the representation of information subjects, and suggest potential paths toward greater inclusion and legitimacy.

In Chapter 2, “How Private Individuals Maintain Privacy and Govern Their Own Health Data Cooperative: MIDATA in Switzerland,” Felix Gille and Effy Vayena explore the Swiss MIDATA cooperative. MIDATA’s members exert cooperative control over the uses of their personal health data through a combination of individual decisions and collective review of project proposals for biomedical research. Within this privacy commons, the board, which reviews research proposals, provides governance and builds trust, while participants across the Swiss population supply the critical resources, namely personal health data.

Chapter 3, “Pooling Mental Health Data with Chatbots,” by Michael Mattioli, presents a critical analysis of applications of conversational agents to treat clinical anxiety. In addition to treating anxiety and depression in real time, these chatbot apps are designed to improve quality of care with time, not only by learning about individual users, but also by creating and using a larger pool of user conversations. Patients who use these chatbots are thus both the source of personal information used as a resource for generating new knowledge and part of the community most directly impacted by its use. Unlike MIDATA, the chatbot governance model does not involve information subject participation, but relies instead on the ethical commitments of its physician creators and patient-informed consent.

In Chapter 4, “Privacy in Practice: A Socio-Technical Integration Research (STIR) Study of Rules-in-Use within Institutional Research,” Chase McCoy and Kyle M. L. Jones study the governance and practice of university data mining and learning analytics using a sociotechnical integration research (STIR) design. Their study probes the value of student data to institutional research, the institutional participants involved with its collection and use, and the ways in which the creation and use of student data knowledge resources are governed. In this case, student information subjects do not participate directly in governance, nor is governance premised on their consent. Instead, privacy governance is based on legal regulation, university policies, and, importantly, collective norms reflecting the ethical commitments of the researchers.

Chapter 5, “Public Facebook Groups for Political Activism,” by Madelyn Sanfilippo and Katherine Strandburg, studies governance of personal information in online social movements that use Facebook as a primary locus for activity. Their empirical study of the Day Without Immigrants movement, the March for Science, and the Women’s March explores the variety of personal information – ranging from personal narratives to contact information – contributed by participants and the complex and polycentric ways in which personal information resources are governed by movement leaders and organizers, informal responses from other participants, and the design of Facebook’s platform. This chapter also serves as a bridge to

the group of studies focused on the ways that privacy governs participation and co-creation of knowledge resources because these movements also must deal with collateral flows of personal information associated with the creation and governance of other types of knowledge resources.

In Chapter 6, “The Republic of Letters and the Origins of Scientific Knowledge Commons,” Michael Madison explores how privacy shaped the historical knowledge sharing practices of “The Republic of Letters,” an early open science regime. The knowledge resources created by this sharing regime were public, both in the sense that they were not secret and in the sense that they were intended to include general, rather than personal, knowledge. Nonetheless, as Madison describes, privacy practices were key to self-organization processes of the Republic of Letters. For example, rules-in-use about personal information sharing both underlay reputational compensation and significantly limited the types of personal information deemed appropriate to share.

In Chapter 7, Brett M. Frischmann, Katherine Haenschen, and Ari Ezra Waldman address “Privacy and Knowledge Production across Contexts.” They compare the rules-in-use governing personal information flows in three distinct contexts: meetings governed by the Chatham House Rule, Gordon Research Conferences, and Broadband Internet Tech Advisory Group (BITAG). Their study shows how these communities use different forms of privacy governance to create trusted environments for information sharing, thereby encouraging participation by diverse contributors to the creation of knowledge resources.

Chapter 8, Scott J. Shackelford’s “Governing the Internet of Everything,” considers the problem of cybersecurity governance in a global Internet system that increasingly involves connected smart devices. He emphasizes the complexity and polycentricity of the cybersecurity governance regime, which involves international, state, commercial, and private actors. Cybersecurity has many aspects, including governance of the ways that various commercial, governmental, and criminal players can exploit users’ personal information. Shackelford warns that the regime complexes addressing cybersecurity may not adequately represent the interests of personal information subjects, particularly those who live in less developed and less powerful states. He argues that the GKC framework and Ostrom’s IAD framework can be used to critically analyze cybersecurity governance in order to develop novel interventions to address these concerns.

In Chapter 9, “Contextual Integrity as a Gauge for Governing Knowledge Commons,” Yan Shvartzshnaider, Madelyn Sanfilippo, and Noah Apthorpe use contextual integrity (CI) as a gauge for evaluating the governance of personal information revealed by users participating in the Internet of Things. Through a survey of public perceptions regarding privacy and IoT devices, they find large gaps between the norms and expectations articulated by some sub-groups of users and the ways that commercial suppliers of smart connected devices govern the aggregation and use of such information. These gaps are evidence that current

4 Madelyn Rose Sanfilippo, Brett M. Frischmann, and Katherine J. Strandburg

governance fails to account for the interests of information subjects. Their study also explores how some smart device users cooperate through user forms to create a distinct knowledge resource of information about how personal information flows in the IoT environment and strategies that users can use to limit the collection of their information, at least to some extent.

Chapter 10, Darakhshan J. Mir's "Designing for the Privacy Commons," examines how the tools and methodologies of design might be used to assess the appropriateness of entrenched norms or rules-in-use associated with privacy. Mir argues that Participatory Design methodology, with its political and ideological commitments to democratic decision-making, may be a particularly promising way to address the deficits in representation of information subjects' interests identified in some cases of personal information governance.

While each of these chapters and case studies is fascinating in its own right, the concluding chapter provides a critical meta-perspective. Taken together, this book's exploration of personal information and its unique connection to information subjects add nuance to our earlier analysis of member-driven, public-driven, and imposed commons governance and bring new themes into focus. Newly salient themes include the role of personal information governance in boundary negotiation and socialization, the potential for conflicts between knowledge contributors and information subjects; the potential adversarial role of commercial infrastructure in imposing commons governance; the role of privacy work-around strategies in responding to those conflicts; the importance of trust; the contestability of commons governance legitimacy; and the co-emergence of contributor communities and knowledge resources. These new studies also confirm and deepen insights into recurring themes identified in previous GKC studies (Frischmann, Madison & Strandburg, 2014; Strandburg, Frischmann & Madison, 2017).

REFERENCES

- Frischmann, Brett M., Michael J. Madison, and Katherine Jo Strandburg, eds. *Governing Knowledge Commons*. Oxford University Press, 2014.
- Nissenbaum, Helen. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.
- Sanfilippo, Madelyn, Brett Frischmann, and Katherine Strandburg, Privacy as commons: case evaluation through the Governing Knowledge Commons framework, *Journal of Information Policy* (8), pp. 116–166 (2018).
- Strandburg, Katherine J., Brett M. Frischmann, and Michael J. Madison, eds. *Governing Medical Knowledge Commons*. Cambridge University Press, 2017.

1

Privacy and Knowledge Commons

Madelyn Rose Sanfilippo¹, Brett M. Frischmann², and Katherine J. Strandburg³

1.1 INTRODUCTION

Although “privacy” and “commons” might on first impression seem conceptually orthogonal or even opposed, a deeper analysis suggests there are insights to be gained from studying information privacy as a question of knowledge commons governance. Privacy often is taken to connote constraint and control over information, while commons often connotes openness and sharing. Neither of these stereotypes, however, are accurate reflections. A more nuanced perspective reveals that sharing and constraint are two sides of the same coin, acting as complements, both in social situations ordinarily conceived in privacy terms and in institutions aimed at creative production through knowledge sharing. Privacy is not simply a matter of constraint, but is more usefully understood, as Helen Nissenbaum has argued, as a matter of “*appropriate flow of personal information*” for specific social contexts (2009, p. 127). When defined as such, it becomes apparent both that privacy is not secrecy and that privacy often involves knowledge sharing. Indeed, true secrecy, in which information is completely unshared (Friedrich, 1971; Neitzke, 2007), is a rarity. Privacy

The original version of this chapter was published in *The Journal of Information Policy* as: Sanfilippo, Frischmann, and Strandburg. “Privacy as Commons: Case Evaluation through the Governing Knowledge Commons Framework.” *Journal of Information Policy*, 8 (2018): 116–166.

¹ Assistant Professor, School of Information Sciences, University of Illinois at Urbana-Champaign; Affiliate Scholar, The Vincent and Elinor Ostrom Workshop in Political Theory and Policy Analysis, Indiana University, Bloomington; Former: Postdoctoral Research Associate, Center for Information Technology Policy, Princeton University and Information Law Institute, New York University; Ph.D., Indiana University, Bloomington; M.I.S., Indiana University, Bloomington; B.S., University of Wisconsin-Madison.

² Charles Widger Endowed University Professor in Law, Business and Economics, Villanova University, Charles Widger School of Law; Affiliated Faculty, The Vincent and Elinor Ostrom Workshop in Political Theory and Policy Analysis, Indiana University Bloomington; Affiliate Scholar, Center for Internet and Society, Stanford Law School; Affiliate of the Princeton Dialogues on AI and Ethics, Princeton University; Trustee, Nexa Center for Internet & Society, Politecnico di Torino. J.D. Georgetown University Law Center; M.S., Columbia University; B.A., Columbia University.

³ Katherine J. Strandburg is the Alfred Engelberg Professor of Law and Director of the Information Law Institute at New York University.

6 *Madelyn Rose Sanfilippo, Brett M. Frischmann, and Katherine J. Strandburg*

ordinarily entails both constraint and flow. Similarly, commons-based knowledge production, at least as understood within the Governing Knowledge Commons (GKC) framework, is rarely free-for-all open sharing, but ordinarily combines sharing practices with constraints to overcome social dilemmas (Frischmann, Madison, and Strandburg, 2014). Thus, privacy may aptly be described not only as contextually appropriate information flow but also as governance of personal information resources.

Given the close affinity between privacy and knowledge commons governance, progress may be made in theoretical and empirical studies of privacy by employing tools developed for the study of knowledge commons governance. In earlier work, Frischmann, Madison, and Strandburg (2014) adapted Elinor Ostrom's Institutional Analysis and Development (IAD) framework for natural resource commons (1990, 2005) to devise a GKC framework for studying commons-based knowledge production. This framework has now been successfully employed in a number of case studies (Frischmann, Madison, and Strandburg, 2014; Strandburg, Frischmann, and Madison, 2017). There is also surprisingly close correspondence between the GKC framework and Nissenbaum's contextual integrity framework for privacy, given their construction for quite different social concerns. Comparing the two, we identify two specific ways in which the knowledge commons approach can help to move the privacy research ball forward.

First, we propose to supplement Nissenbaum's conceptions of "transmission principles" and "context-relevant information norms" with the more politically and procedurally grounded conceptions of rules-in-use and governance employed in commons studies. In Nissenbaum's framework, appropriate flows of information are distinguished, in the first instance, by compliance with "transmission principles," defined as "terms and conditions under which such transfers ought (or ought not) to occur" (Nissenbaum, 2009, p. 145) between specific parties in a specific context. The "transmission principles" observed in a specific situation are examples of what Ostrom called "rules-in-use." Ostrom's concept of "rules-in-use" differentiates between nominal rules "on the book" and the actual (and perhaps unanticipated) practices that emerge from interactions within often complex structures of formal and informal institutional arrangements. Ostrom further taxonomized "rules-in-use" into an "institutional grammar" that encompasses rules, social norms, and strategies (Crawford and Ostrom, 1995), as well as individual tactics of compliance and avoidance, power dynamics, and enforcement mechanisms. This approach can be used to add depth to our understanding of the privacy transmission principles observed in various real-world situations. The "rules-in-use" concept allows sweeps beyond information transmission to include the possibility of other sorts of constraints, such as rules-in-use governing how personal information may appropriately be exploited.

Under Nissenbaum's framework, when transmission principles are contested, eroded, or changed as a result of social and technological changes, their normative

validity is tested against “context-relevant informational norms” and overarching ethical principles. The origins of contextual norms governing appropriate information flow are exogenous to Nissenbaum’s analysis. The commons governance perspective encourages us to look behind the curtain to investigate the *origins* and dynamic characters of both nominal rules and rules-in-use, and to interrogate the potentially contested legitimacy of the formal and informal processes that produce them. We believe that issues of procedural legitimacy and distinctions between nominal rules and rules-in-use are central both to descriptive understanding of privacy and to normative evaluation and policymaking. Governance and legitimacy may be particularly important for the most perplexing privacy issues, which often involve overlapping ethical contexts or contested values.

Second, we propose the knowledge commons framework as a rigorous, yet flexible, means to systematize descriptive empirical case studies of real-world contexts; it is primarily an explanatory approach, rather than a descriptive theory, and structures analysis of nested and networked policy instruments and management strategies (Bennett and Raab, 2006). Accurate empirical understanding is an essential basis for constructing and evaluating theory and for effective policy design. Privacy, understood as “appropriate” personal information flow, takes complex and variable forms that can only be understood by delving deeply into specific real-world situations. If general principles are to be gleaned from case studies of such various and heterogeneous situations, a systematic framework is needed. The IAD framework was applied successfully by Ostrom and collaborators to derive general “design principles” from case studies of natural resource commons (Ostrom, 1990). The accumulation of case studies employing the IAD-derived GKC framework is at an earlier stage, but general insights and testable hypotheses have already started to emerge (Frischmann, Madison, and Strandburg, 2014; Strandburg, Frischmann, and Madison, 2017). We anticipate that using the enhanced GKC framework proposed here to structure systematic case studies of how personal information flows are governed in various real-world contexts will lead to similar progress in our understanding of privacy.

This chapter aims to convince readers that the commons approach to information privacy has a good chance of producing new and useful insights. We thus supplement our conceptual discussion of the approach with a demonstration study in which we identify and analyze privacy issues that were implicit in previously studied knowledge commons cases. Those studies have produced insights into a variety of aspects of knowledge production within communities, ranging from the various social dilemmas communities may face when seeking to achieve their objectives to the institutional governance choices they rely on to overcome those dilemmas. A previous analysis of knowledge-sharing regimes elucidated differences along four distinct community designs: centralized, intermediate distributed, fully distributed, and noncommons (Contreras and Reichman, 2015). Similarly, our meta-analysis, focusing on personal information sharing, uncovered three distinctive

8 *Madelyn Rose Sanfilippo, Brett M. Frischmann, and Katherine J. Strandburg*

patterns of rules-in-use based on whether the governance was public driven, member driven, or imposed by leadership or a platform. This reanalysis of previous case studies is intended to be exemplary, rather than representative of the range of situations in which privacy debates arise, so it is likely that additional patterns will emerge from case studies undertaken with privacy in mind. Nevertheless, the meta-analysis presented here uncovers interesting empirical patterns and raises issues that are worthy of further exploration; in particular, the knowledge commons perspective highlights the interdependence between knowledge flows aimed at creative production and personal information flows. It also demonstrates that a contextualized understanding of privacy requires a broad conception of “personal information” that extends well beyond information that is ordinarily deemed “sensitive.” For example, inappropriate flows of information such as an individual’s views, opinions, or ideas can stifle socially valuable information sharing or have other undesirable effects.

This meta-analysis demonstrates that those who systematically study knowledge commons governance with an eye toward knowledge production routinely encounter privacy concerns and values, along with rules-in-use that govern appropriate personal information flow. In the same way, we anticipate that many, if not most, communities within which privacy is a hotly contested issue are also dealing with corresponding questions about knowledge production, sharing, curation, and use – or more generally, knowledge governance. In sum, while this chapter does not attempt a new conceptualization of privacy per se, it contends that institutional analysis can be an important conceptual and empirical aid to privacy research and that understanding privacy as governance of personal information flows can illuminate otherwise underappreciated facets of knowledge commons arrangements.

1.2 THEORETICAL BACKGROUND

In order to explore the utility of integrating the GKC framework (1.2.1) with Nissenbaum’s Contextual Integrity framework (1.2.2), it is first necessary to understand and compare them, and to identify points of synergy and possibilities for augmentation (1.2.3), including research questions to be explored in further developing the GKC framework.

1.2.1 *The GKC Framework*

Commons governance of natural resources is often explored through Ostrom’s IAD framework. Commons, as used in the literature upon which we build here, refers to community management or governance of resources. “The basic characteristic that distinguishes commons from non-commons is *institutionalized sharing of resources* among members of a community” (Madison, Frischmann, and Strandburg, 2009, p. 841). Commons governance can take many forms and need not involve the kind of

complete openness often associated with discussions of “the commons” or “the public domain” in the legal literature, nor should it be conflated with the type of resources that are managed.

Ostrom’s work initially emphasized the appropriateness of commons governance for “common pool resources,” meaning “a natural or man-made resource system that is sufficiently large as to make it costly (but not impossible) to exclude potential beneficiaries from obtaining benefits from its use” (Ostrom, 2015, p. 4). In economic terms, common pool resources are rivalrous and nonexcludable and commons governance of such resources generally aims to address so-called tragedies of the commons, social dilemmas associated with overuse – congestion, depletion, and destruction. Commons governance is used by a wide variety of communities to manage many different types of resources, however, and responds to various obstacles to sustainable sharing and cooperation. Some of those obstacles derive from the nature of the resources and others derive from other factors, such as the nature of the community or external influences.

When we refer to knowledge commons, we mean commons governance applied to knowledge resources, broadly defined, where:

Knowledge refers to a broad set of intellectual and cultural resources. ... We emphasize that we cast a wide net and that we group information, science, knowledge, creative works, data, and so on together. (Frischmann, Madison, and Strandburg, 2014, p. 2)

In this sense, knowledge resources may lie at any point along the data, information, knowledge, and wisdom hierarchy (Henry, 1974). Personal information, broadly defined, is one type of knowledge resource, which can produce value when it is shared and managed appropriately.

As recognized by Hess and Ostrom (2007) and confirmed by later GKC studies, “sharing of knowledge often is sustained by commons governance.” Indeed, case studies of knowledge commons have illustrated the use of commons governance to manage not only knowledge, which is a classic public good,⁴ but also classic private goods, such as money, that must be shared to accomplish a community’s goals and objectives.

We anticipate that commons governance will often be applied to flows of personal information for related, but somewhat distinct reasons. If personal information can flow without constraint, the subjects of the information may either be disinclined to share it at all, opting for secrecy, or, if secrecy is not possible, may be unfairly harmed by the flow. Commons governance can provide for the beneficial and managed flow of personal information within a legitimate and trusted institutional structure, thus encouraging subjects to share it in a particular social setting and reducing the potential that harm will result from doing so.

The GKC framework (which is adapted for knowledge resources from Ostrom’s IAD framework) is represented in Figure 1.1.

⁴ More extensive discussions of the public goods nature of knowledge are presented by Frischmann, Madison, and Strandburg (2014, introduction, p. ix) and Ostrom and Hess (2007).

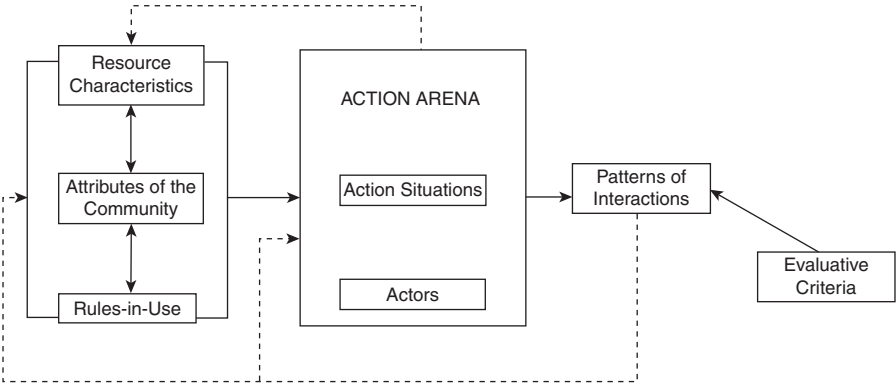


FIGURE 1.1 The GKC framework

Using the IAD framework, Ostrom and colleagues explored patterns of community interactions (McGinnis, 2011). *Action arenas* serve as the core units of IAD and GKC analysis, functioning as policy analysis equivalent of social action and interaction settings (Burns and Flam, 1987 or Goffman’s frames, 1974). An action arena is simply a recurring type of situation in which community actors interact with one another. Interactions in an action arena produce outcomes, denoted here as patterns of interaction, which can then be evaluated according to some community or socially generated criteria. The figure depicts how effects flow between conceptual building blocks. Thus, resource characteristics, community attributes (including members and roles), and a set of governing “rules-in-use” are inputs to an action arena. Patterns of interactions accumulate, feeding back to create new action situations and influencing resource characteristics, community attributes, and rules-in-use. Knowledge resources are often produced and defined by the community. The knowledge outputs of some knowledge commons action arenas must themselves be managed by the community and may be inputs to further knowledge production. This feedback, between a community’s activity and its available knowledge resources, justifies community-level analysis, emphasizing questions related to group interactions and outcomes, rather than user-level analysis, emphasizing questions about individual experiences.

Focusing on action arenas facilitates examination of resource sharing in dynamic local contexts, as opposed to simply examining interactions in broad contexts (Ostrom, 2005). The “action arena” concept is flexible and can be applied at a variety of levels of generality, depending upon the question of interest to the analyst. Analyzing an action arena is meaningful only if one can specify resource characteristics, community attributes, and rules-in-use that are “exogenous” or fixed over a number of action situations and if one can describe