

# **Privacy in Context**

## **Technology, Policy, and the Integrity of Social Life**

Helen Nissenbaum

<http://www.nyu.edu/projects/nissenbaum/index.html>

The Table of Contents and Introduction to Professor Nissenbaum's **Privacy in Context: Technology, Policy, and the Integrity of Social Life** are attached. The book is available from Stanford University Press. For more information, please see <http://www.sup.org/book.cgi?id=8862>.

# PRIVACY IN CONTEXT

*Technology, Policy, and  
the Integrity of Social Life*

Helen Nissenbaum

Stanford Law Books  
An Imprint of Stanford University Press  
Stanford, California

--1  
—0  
—+1

Stanford University Press  
Stanford, California

© 2010 by the Board of Trustees of the Leland Stanford Junior University. All rights reserved.

No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or in any information storage or retrieval system without the prior written permission of Stanford University Press.

Printed in the United States of America on acid-free, archival-quality paper

Library of Congress Cataloging-in-Publication Data  
Nissenbaum, Helen Fay.

Privacy in context : technology, policy, and the integrity of social life / Helen Nissenbaum.

p. cm.

Includes bibliographical references and index.

ISBN 978-0-8047-5236-7 (cloth : alk. paper)—ISBN 978-0-8047-5237-4 (pbk. : alk. paper)

1. Privacy, Right of—United States. 2. Information technology—Social aspects—United States. 3. Information policy—United States. 4. Social norms. I. Title.

JC596.2.U5N57 2010

323.44'80973—dc22 2009026320

Typeset by Westchester Book Composition in Minion, 10/14

-1—  
0—  
+1—

## Contents

Acknowledgments	ix
Introduction	1
PART I: Information Technology's Power and Threat	
1 Keeping Track and Watching over Us	21
2 Knowing Us Better than We Know Ourselves: Massive and Deep Databases	36
3 Capacity to Spread and Find Everything, Everywhere	51
PART II: Critical Survey of Predominant Approaches to Privacy	
4 Locating the Value in Privacy	67
5 Privacy in Private	89
6 Puzzles, Paradoxes, and Privacy in Public	103

--1  
—0  
—+i

PART III: The Framework of Contextual Integrity

7	Contexts, Informational Norms, Actors, Attributes, and Transmission Principles	129
8	Breaking Rules for Good	158
9	Privacy Rights in Context: Applying the Framework	186
	Conclusion	231
	Notes	245
	References	257
	Index	281

## Introduction

INFORMATION TECHNOLOGY IS CONSIDERED A MAJOR THREAT TO privacy because it enables pervasive surveillance, massive databases, and lightning-speed distribution of information across the globe. In fact, privacy has been one of the most enduring social issues associated with digital electronic information technologies. A fixture in public discourse at least since the 1960s, when the dominant concern was massive databases of government and other large institutions housed in large stand-alone computers, concerns have multiplied in type and extent as radical transformations of the technology have yielded the remarkable range of present-day systems, including distributed networking; the World Wide Web; mobile devices; video, audio, and biometric surveillance; global positioning; ubiquitous computing; social networks; sensor networks; databases of compiled information; data mining; and more. Associated with each of these developments is a set of worries about privacy. Whether expressed in the resigned grumbles of individuals, the vocal protests of advocacy groups and eloquent politicians, or the pages of scholarly publications and popular media, the common worry time and again is that an important value is a casualty of progress driven by technologies of information.

Countless books, articles, and commentaries call for reform in law and policy to shore up defenses against the erosion of privacy due to swelling ranks of technology-based systems practices. Many of them argue that protecting privacy means strictly limiting access to personal information or

--1  
—0  
--+1

assuring people's right to control information about themselves. I disagree. What people care most about is not simply *restricting* the flow of information but ensuring that it flows *appropriately*, and an account of appropriate flow is given here through the framework of contextual integrity. The framework of contextual integrity provides a rigorous, substantive account of factors determining when people will perceive new information technologies and systems as threats to privacy; it not only predicts how people will react to such systems but also formulates an approach to evaluating these systems and prescribing legitimate responses to them.

Almost as many who have taken up the subject of privacy in relation to information technology have declared it deeply problematic, referring not only to questions and disagreements about its value, benefits, and harms but to its conceptual morass. Attempts to define it have been notoriously controversial and have been accused of vagueness and internal inconsistency—of being overly inclusive, excessively narrow, or insufficiently distinct from other value concepts. Believing conceptual murkiness to be a key obstacle to resolving problems, many have embarked on the treacherous path of defining privacy. As a prelude to addressing crucial substantive questions, they have sought to establish whether privacy is a claim, a right, an interest, a value, a preference, or merely a state of existence. They have defended accounts of privacy as a descriptive concept, a normative concept, a legal concept, or all three. They have taken positions on whether privacy applies only to information, to actions and decisions (the so-called constitutional rights to privacy), to special seclusion, or to all three. They have declared privacy relevant to all information, or only to a rarefied subset of personal, sensitive, or intimate information, and they have disagreed over whether it is a right to control and limit access or merely a measure of the degree of access others have to us and to information about us. They have posited links between privacy and anonymity, privacy and secrecy, privacy and confidentiality, and privacy and solitude.

Believing that one must define or provide an account of privacy before one can systematically address critical challenges can thwart further progress. Those who hold that a credible account is one that maps natural usage are confronted with a fractured, ambiguous, perhaps even incoherent concept and are understandably hard-pressed to unify the disparate strands of meaning. Maintaining all these meanings while delineating a concept to support policy, moral judgment, and technical design seems a hopeless ambition.<sup>1</sup>

-1—  
o—  
+1—

Those who recognize the perils of inclusiveness attempt to purify the concept by trimming away some of the inconsistency and ambiguity, declaring certain uses wrong or confused. This has meant disputing the proper application of privacy to so-called constitutional cases, or it has meant rejecting control over information as part of the meaning of privacy in favor of degree of access, or vice versa.<sup>2</sup> A third strategy is to stipulate a precise definition necessary for a specific application without necessarily connecting this with natural etymology or a full natural meaning of the term; this is common in the works of computer scientists and engineers and necessary in relation to the purposes they clearly specify.<sup>3</sup>

In contrast, this book does not mediate its investigation of the unsettling stream of systems and practices through the concept of privacy. It does not carve a pathway through the conceptual quagmire to claim a definition—its definition—of privacy. Nevertheless, it is a book about privacy because it explains why the huge and growing set of technical systems and technology-based practices have provoked and continue to provoke anxiety, protest, and resistance in the name of privacy.

The framework of contextual integrity identifies the roots of bewilderment, resistance, and sometimes resignation expressed by experts and non-experts alike. According to the framework, finely calibrated systems of social norms, or rules, govern the flow of personal information in distinct social contexts (e.g., education, health care, and politics). These norms, which I call context-relative informational norms, define and sustain essential activities and key relationships and interests, protect people and groups against harm, and balance the distribution of power. Responsive to historical, cultural, and even geographic contingencies, informational norms evolve over time in distinct patterns from society to society. Information technologies alarm us when they flout these informational norms—when, in the words of the framework, they violate contextual integrity.

As troubled as we might be by technologies that diminish control over information about ourselves, even more deeply troubling are those that disregard entrenched norms because, as such, they threaten disruption to the very fabric of social life. To be sure, not all systems that alter the flow of information are cause for alarm, for there are clear cases of new information devices and systems that serve societal as well as context-based values, ends, and purposes better than those we already have in place (e.g., promoting intellectual development, health and well-being, and vibrant democracy). In such

—1  
—0  
—+1



cases, the systems in question generally are and should be accepted, even celebrated.

### Privacy and Personal Information Flow

Privacy is the initial organizing principle defining the scope of this book because, historically, it has been the term in which concerns, anxieties, and protests have been expressed. As the book proceeds, however, it frames its claims in terms of personal information flows, not only less encumbered with normative assumptions but useful for characterizing fundamental similarities at the heart of an otherwise disparate array of systems and devices. Because the book also seeks to provide an evaluation of these systems and devices in moral and political terms, the language of information flow allows us to sidestep certain of the disagreements and confusion associated with the concept of privacy and avoid potential question begging without sacrificing precision.

A related point about terminology: there is great ambiguity in the way “personal information” is used. Colloquially and in contexts of privacy law and policy, as well as academic research, it can mean sensitive or intimate information, any information about a person, or only personally identifying information. Here and throughout the book, following usage practices in the policy community, I use it to mean information about an identifiable person—for example, as defined in the European Union Directive, “personal data shall mean any information relating to an identified or identifiable natural person (‘data subject’); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity.”<sup>4</sup>

### Technology and the Socio-technical

In the study of technology as a social and political phenomenon, many works acknowledge its diverse meanings but seek to build on those that ring true and are theoretically useful. In this book, too, it is important to explain, briefly, what I mean by technology as well as the related notion of a socio-technical system. To begin, consider the familiar telephone sitting on your desk. Upon initial reckoning you could see it as a self-standing technical device (think of the phone in the box when you purchased it), but its capacity to function as a

-1—  
0—  
+1—

telephone, enabling communication at a distance, requires that it be connected to a complex telecommunications system including all necessary hardware and software. Beyond these, a proper functioning telecommunications system depends on a host of social, political, and economic arrangements. Because complex interdependencies such as these are integral to the functioning of almost all of the technologies of contemporary societies, it is misleading to think of the study of technology's impacts on people and societies as an investigation of a stand-alone physical device—wires, hardware, and software. Rather, the object of study is the device construed in terms of key social interdependencies, as responsible for its features, function, and impact as for its physical characteristics.

In the case of the telephone, its observable characteristics in the box but disconnected from broader technical and social systems might include, say, the sound tones it makes when keys are pressed, its capacity to cause a gash on someone's head if dropped from a certain height, or its aesthetic and ergonomic qualities. But many of the most interesting features of "the telephone," including past and predicted societal impacts, are due to its properties as an embedded device. When observing, say, its effects on workplace hierarchy, on the home, on friendship, on the aged, on law enforcement, on urban development, and so forth, we do not mean the telephone in the box but the telephone connected to a telecommunications system, regulated by a host of technical standards, public policies, and even social and cultural norms. These dependencies are more evident in systems and devices that require customization, such as a closed-circuit television (CCTV) security system requiring knowledge of movement patterns in the area targeted for surveillance, but they are evident in many others, such as the automobile, and even in those we consider "plug and play," such as dishwashers and televisions.

Conscious of these complex interdependencies when referring to "the telephone," "the automobile," "the computer," and "the Internet," and wishing to highlight them, scholars of the social and humanistic study of technology refer to them as socio-technical devices and systems. Likewise, it is important to bear in mind that the devices and systems—technologies—of concern in this book, namely, those altering the flow of personal information in radical and radically worrying ways, are socio-technical. For example, radio frequency identification (RFID) technology or vehicle safety communications systems (VSCS) (discussed in Chapter 1), which might at first glance appear to be *pure* technologies, cannot be properly understood without grasping

—1  
—0  
—+1

their definitional components, “identification,” “safety,” and “communication,” all thoroughly social. Thus, when investigating why and how technological devices and systems, including RFID technologies, provoke anxiety, protest, concern, and resistance in the name of privacy, I am thinking of them as *socio-technical*; they affect us not purely by dint of physical or material properties but by properties they acquire as systems and devices embedded in larger material and social networks and webs of meaning. Accordingly, the terms “socio-technical systems and devices” and “technology-based systems and practices” are used throughout the book, but even when “technology” is used alone it should be read with awareness of the larger picture.

#### Contextual Integrity as a Justificatory Framework

The starting place of this book is the myriad socio-technical systems, devices, and associated practices that control, manage, and steer the flow of personal information, particularly those that have precipitated radical changes, aroused suspicion, caused anxiety, and drawn protest and resistance. They are experienced and registered as threats to and violations of privacy not only individually, case by case, but in aggregate amounting to a social crisis, a watershed: privacy itself is in jeopardy not merely in one or another instance but under attack as a general, societal value. The primary mission of this book is to confront and give a moral and political account of this pileup of technologies and practices, to pinpoint and understand sources of concern, and to provide a framework for expressing and justifying constraints expressed as social norms, policies, law, and technical design.

It is important to recognize, however, that reactions to these systems are not uniform. Nor is it the case that all systems affecting the flows of information are resisted; some are not only ignored and tolerated but are even welcomed and celebrated. In healthcare environments such as hospitals and nursing homes, for example, a plethora of devices such as blood-pressure monitors, pulse oximeters, ECGs, and EEGs complement attentive, responsive caregivers and enhance the close monitoring and recording of patients' condition that is one of the hallmarks of high-quality care. By contrast, video surveillance of public parks, frequent shoppers' card store loyalty programs, government wiretapping of telephone calls, and monitoring online transactions, also resulting in alterations of information flow, are greeted with suspicion and resentment. A satisfactory moral and political account needs to explore and explain these

-1—  
0—  
+1—

contrasts, to understand the sources of disagreement and conflict, and to offer approaches to resolving or at least meliorating them.

As the privacy conundrum has grown in public awareness it has attracted the attention of leaders in all social sectors, including business, government, and education, as well as scholars and researchers across the disciplines. Respondents have taken on its challenges in various ways, advocating certain public policies or promulgating guidelines within business, financial, and healthcare organizations. Legal scholars have developed and championed approaches to privacy law, both advocating for certain regulations as well as recommending to courts how to interpret existing law and past cases to afford adequate protection of privacy rights in conflicts and disagreements over the flows of personal information. Concern over privacy has also reached the scientific world of technical development and deployment, not only yielding a dedicated array of privacy preserving technologies but also leading to the adoption of hardware and software design standards by companies and consortia.

The framework of contextual integrity developed in this book does not lie fully within any one of these efforts, though it complements them (and vice versa). Like them, it attends to policy and regulation, court decisions and law, and technology design and implementation, and it prescribes or expresses support for certain directions over others. Its primary mission, however, is to articulate a foundation for these directions so we may answer questions not only of the form: *what* policies, *what* court decisions, *what* technical standards and design features, but *why* these, with answers rooted in humanistic moral and political traditions of contemporary liberal democracies. The framework provides a way to characterize systems and practices dramatically affecting the flows of information. It provides a language, a form of expression, for explaining when and why they are troubling, whether the balance of reasons favors one side or another, and (in cases of conflict) serves as the basis for prescribing courses of action, decisions, policies, and designs. In relation to perennially hard cases, the framework of contextual integrity enriches our expressive capacity for adjudication.

#### Comparing Contextual Integrity with Other “Justificatory Approaches”

There is, by now, a huge body of work on privacy threats of technologies developed by academic researchers, public interest advocates, legal practitioners

—-1  
—0  
—+1

and theorists, technology designers, and policy makers. Their purposes are not only to describe and prescribe but, like mine, to articulate systems of reasoning, to articulate justificatory frameworks—though not necessarily described in these terms by their authors.

One such approach, frequently adopted in policy-making, legal, and advocacy arenas, highlights the interest politics inherent in controversial systems and practices. Interested parties and their advocates scrutinize these systems for potential impacts on respective rights, interests, benefits, and harms. In general, controversial systems are ones found to be unbalanced in the interests they serve. Uncontroversial acceptance of healthcare monitoring systems can be explained by pointing to the roughly even service to the interests of patients, hospitals, healthcare professionals, and so on, while video surveillance in public parks is perceived to serve the interests of the watchers (e.g., law enforcement personnel) while diminishing the liberties of parkgoers, and online logging and surveillance is seen as promoting the interests of advertisers and marketers but diminishing consumer bargaining power and autonomy.

It is not uncommon for the resolution of such conflicts to involve hard-fought interest brawls, each side campaigning on behalf of privacy or against it, in favor of regulation of practice or the opposite (so-called self-regulation), for more invasive monitoring or less, and so on. Business and government interests in accumulating and using personal information have often prevailed in the face of public complaints, with a few well-known exceptions. One such exception, Lotus Marketplace: Households, a consumer data aggregation on millions of American households that was to have been distributed and sold in CD-ROM format, was quashed by its corporate backers in 1991 as a consequence of public outcry and well-orchestrated resistance by privacy advocacy organizations. In another, also following public, media, and internal criticism, Congress cut funding for the Defense Advanced Research Projects Agency's Office of Information, which would have administered a counterterrorism program proposed by Admiral John Pointdexter (formerly a U.S. National Security Advisor to Ronald Reagan) using "total information awareness" to pre-empt future attacks.

The trouble with settling conflicts through brute clashes among interest holders is the advantage it gives to those possessing advantages of power, resources, and the capacity for unremitting persistence, favoring corporate and governmental actors over the public interest in the long run. In

-1—  
0—  
+1—

retrospect, it is clear that the victories of Lotus Marketplace: Households and Total Information Awareness (TIA, later dubbed “Terrorist Information Awareness”) have been short-lived as they have been resurrected in other more potent, more insidious forms, namely, private sector information service providers such as ChoicePoint (discussed in Chapter 2) as well as fusion centers, creations of state and city government for sharing information among agencies, and data aggregates developed by national security agencies. Although there may be several reasons why Lotus Marketplace Households and TIA failed and fusion centers and contemporary information brokers continue to flourish, still sorely lacking from public discussions of these systems and programs is a clear understanding of what makes one acceptable and another unacceptable. Still missing, in other words, is a justificatory platform or framework to reason in moral terms about them. A brute competition among interests might win the day, but for a particular case to serve as a precedent that carries forward into the future, advocates need to be able to show that a framework of widely accepted principles supporting the first case can also apply to cases in question. When corporate backers of Lotus Marketplace Households capitulated, they acknowledged no moral (or legal) wrongdoing, saying simply that it had become a public relations nightmare. In bowing to public outcry but conceding no ground, in principle they denied their critics the power of precedence and entrenched the interest brawl as a salient form of settling privacy disputes.

Although interest politics may be disguised in sophisticated rhetoric when concerned parties attempt to link the interests of others, even those of the public, with their own, other approaches make the case for privacy explicitly in terms of universal human principles and values. Countless works, many of them brilliant, have defended privacy as a fundamental human right (not merely a preference or an interest) by linking it to other values with long-standing moral, political, and legal pedigrees. These works have shown privacy to be a form and expression of self-ownership, an aspect of the right to be let alone, a cornerstone of liberty and autonomy, or a necessary condition for trust, friendship, creativity, and moral autonomy. The shortcoming of these works, and this approach, is not that it gets things wrong, generally speaking, but that it leaves a gap. This gap is acutely felt for those who are interested in analyzing controversial systems and in the practical mission of prescribing sound decisions in relation to them.

—1  
—0  
—+1

So, where does the framework of contextual integrity fit? A spatial metaphor may help answer this question. Consider a few of the controversial questions confronting us at the time of writing this book: whether it is morally wrong for Google Maps' Street View to include images of identifiable individuals (or their possessions) without permission, whether the FBI should be allowed to coerce librarians to divulge a library's lending logs, whether Internet service providers are entitled to track customers' clickstreams and sell them at will, whether one may post a tagged group photograph of others on one's Facebook page, whether insurance companies violate client privacy when they generate massive databases pooled from information about their clients, whether the police should be permitted to erect covert license plate recognition systems at public intersections, and so on.

Drawing on the spatial metaphor, let us place interest politics on the bottom, on the hard ground of concrete, gritty, detail. Whether the interest-brawls, as I have called them, are won and lost through force of rhetoric, brute resources, or the give-and-take of compromise and accommodation, reached by parties themselves or imposed by third-party mediators such as the courts or the marketplace, they involve concrete detail specific to respective cases in question.

If interest-brawls are conceived as taking place at ground level, appeals to universal human values and moral and political principles take place in the stratospheres of abstraction because resolutions for real-world disputes are sought in the realms of general values and principles. In the case of Street View, one may argue that it violates self-ownership; the FBI may be accused of overstepping principles of liberal democracy constituting the relationship between citizens and government actors; insurance companies might be accused of undermining personal autonomy. Although insight can be gained in identifying connections to higher-order values and principles, a common challenge to such reasoning is conflict, not only at the ground level of interests but at vaunted levels of abstraction—for example, in noting that individual liberty conflicts with national security, personal autonomy with freedoms of business institutions implicit in a free-market economy, and moral autonomy with social order.

Between the ground and the heavens, according to the picture I am imagining, is the realm of the social, and it is in this realm that contextual integrity fits. This middle realm holds a key to explaining why people react to real-world disputes in the ways they do and why they frequently express their

-1—  
 0—  
 +1—

alarm in terms of the erosion of privacy. Although it remains crucial to the understanding of these disputes that we grasp the configurations of interests, values, and principles present in them, our capacity to explain them is diminished if we attend to these elements alone, blind to the sway of social structures and norms. Tethered to fundamental moral and political principles, enriched by key social elements, the framework of contextual integrity is sufficiently expressive to model peoples' reactions to troubling technology-based systems and practices as well as to formulate normative guidelines for policy, action, and design.

### Book Outline

The book comprises three parts, each comprising three chapters. Part I is devoted to technology, Part II to predominant approaches to privacy that have influenced and informed contextual integrity, Part III develops the framework of contextual integrity, circling back to technologies discussed in Part I and illustrating its application to these and others.

#### *Part I*

Part I is a contemporary snapshot of the landscape of technologies and socio-technical systems, including a few detailed close-ups. One at a time, it is not difficult to recognize a technical system or technology-based practice as one that threatens privacy. But when considered together in a single class, they present an array of bewildering variety and the task of classifying them according to common features proves to be daunting. Yes, they affect the flows of personal information and threaten privacy, but that is not a terribly illuminating observation. Given its importance to the project as a whole, finding a satisfactory way of characterizing these technology-based systems and practices was, however, a challenge that could not be finessed.

The structure on which I settled is described in the three chapters, each mapping onto one of three capacities: (1) tracking and monitoring, (2) aggregation and analysis, and (3) dissemination and publication. It is important to note that systems and practices do not fit uniquely into only one of these categories but may incorporate more than one of the capacities and possibly even all three.

Chapter 1 surveys the vast array of technology-based systems and practices whose capacity to track and monitor people lies at the root of privacy

—1  
—0  
—+1



worries, protests, and resistance. Expansion of this array is due not to any single technological breakthrough but rather to many breakthroughs, amplified by incremental advances in supporting technologies, such as input or information capture devices (such as digital photography and sound recording), digital encoding algorithms, network transmission mechanisms, information storage capacity, and general software controls. These, in various combinations and permutations, constitute the substrate for monitoring and tracking. Although attention has focused primarily on highly visible applications such as video surveillance, wiretapping, and online monitoring of Web transactions, this category includes a slew of less obtrusive, more specialized systems, some already in operation and many others under development and poised to enter the mainstream. From the mundane frequent shoppers' card to the myriad services cropping up in all walks of life, such as the one offered by wireless telephone providers to parents to track their children's movements (e.g., Verizon's "Chaperone"), to the well-meaning "intelligent homes" equipped with sundry embedded sensors enabling the elderly to live independently, to the somewhat more sinister watchfulness of workplace e-mail surveillance, these systems keep track episodically or continuously of people's whereabouts, activities, and attributes. The chapter offers a selective survey of this class of systems, with an in-depth focus on radio frequency identification technology.

Systems for monitoring and tracking are often highly visible, but the legendary powers of computing and information technologies to store and manipulate information have, arguably, contributed far more to technology-based privacy challenges. Chapter 2 presents a small sample of the breathtaking array of systems facilitated by these capacities, such as "back-end" storage capacities that are essential for almost all tracking and monitoring systems. Increasingly effective scientific approaches to organizing, analyzing, manipulating, storing, retrieving, and transmitting information make the information stored in stand-alone and distributed databases increasingly useful. Furthermore, these competencies, once affordable only to government and large financial institutions, are now widely dispersed, with the result that information is lodged not only in obvious and familiar places but in places we cannot even begin to guess. Drawing on media reports and scholarship, Chapter 2 surveys some of these, focusing on the emergent niche of corporate information brokers, such as ChoicePoint, providing a wide range of information services in the private and public sectors.

-1—  
0—  
+1—

A third general capacity afforded by computing and information sciences and technologies is the capacity to disseminate, transmit, communicate, broadcast, or publish and publicize information. As with the other two, this capacity was initially exploited by a few large institutional actors such as news and other centralized broadcast media but rapidly has promulgated throughout society, due mostly to progress in and wide adoption of digital communications networks, predominantly the Internet and the World Wide Web. Chapter 3 discusses hard privacy issues raised by these remarkable technology-based changes, with extended attention devoted to two cases: (1) the placement of public records, including court records, on the Web and (2) the intriguing challenges posed by so-called Web 2.0 applications, including social networking sites such as Facebook and MySpace.

#### *Part II*

Part II offers a critical survey of predominant approaches to privacy, sampling explicit principles guiding law and policy as well as several leading theoretical contributions. It is impossible, in the scope of this book, to provide detailed and systematic accounts of individual theories. Rather, my intention is to explore predominant themes and principles as well as a few of the well-known theories that embody them.

Chapter 4 identifies two general approaches to explaining the sources of privacy's importance as a right, or a value, deserving moral consideration as well as legal protection. One attributes the value of privacy to the crucial role it plays in supporting other moral and political rights and values. The other locates privacy's value in the critical role it plays protecting the sphere of the private. Most of Chapter 4 is devoted to the first of these.

One of the charges frequently leveled against privacy advocacy and scholarship is that its sprawling domain is difficult, if not impossible, to capture with a coherent and distinctive concept. Chapter 5 discusses the approach to containing this conceptual sprawl, based on the private-public dichotomy, which holds that a right to privacy extends only across zones of life considered private. A commitment to this thesis, which has been compelling to both practitioners and scholars, is evident in the literature on privacy as well as in policy formation and key court rulings where privacy protection has attached to private information, private space, and private activities but not to their public counterparts.

—1  
—0  
—+1

Chapter 6, constructing a bridge to Part III, highlights challenges posed by technology-based systems to these theories and paradigms. One recurring skeptical challenge, for instance, cites the lack of concern many people seem to demonstrate in day-to-day behaviors, contradicting claims that privacy is a deeply important moral and political value that deserves stringent protection. Another is the clearly evident cultural and historical variation in commitments to privacy, hard to explain if privacy is supposed to be a fundamental human right. A third points to the difficulty of resolving conflicts between privacy and other moral and political values, such as property, accountability, and security. Most puzzling of all, however, is the problem of privacy in public, which challenges accounts of privacy that rely on the private-public dichotomy. The framework of contextual integrity is able to respond to them all.

### *Part III*

Part III explicates the framework of contextual integrity. The central claim is that contextual integrity captures the meaning of privacy in relation to personal information; predicts people's reactions to new technologies because it captures what we care about when we question, protest, and resist them; and finally, offers a way to carefully evaluate these disruptive technologies. In addition, the framework yields practical, step-by-step guidelines for evaluating systems in question, which it calls the CI Decision Heuristic and the Augmented CI Decision Heuristic.

Chapter 7 introduces key features of the framework beginning with the basic building block of social contexts; the underlying thesis is that social activity occurs in contexts and is governed by context-relative norms. Among these, informational norms govern the flow of information about a subject from one party to another, taking account of the capacities (or roles) in which the parties act, the types of information, and the principles under which this information is transmitted among the parties. We can think of contextual integrity as a metric, preserved when informational norms within a context are respected and violated when they are contravened. Whether contextual integrity is preserved or violated by a newly introduced system or practice is claimed to be predictive of people's reactions—whether they protest, accept, or even welcome it.

Chapter 8 addresses a potential limitation of the framework of contextual integrity, which, to this point, requires compliance with entrenched social

-1—  
0—  
+1—

norms. To avoid the charge of stodginess and conservatism, it needs to incorporate ways not only to detect whether practices run afoul of entrenched norms but to allow that divergent practices may at times be “better” than those prescribed by existing norms. This requirement is accommodated by an augmented analysis that begins with a presumption in favor of entrenched or normative practices, based on the belief that they are likely to reflect settled accommodation among diverse claims and interests. A presumption in favor does not, however, preclude legitimate challenges, and the approach developed in Chapter 8 looks to a context’s internal purposes, ends, and values for benchmarks against which entrenched and novel practices may be evaluated and compared. Accordingly, the augmented framework of contextual integrity tells us that new technologies deserve to be embraced when they help achieve important social and context-based ends more effectively than was possible prior to their use.

In Chapter 9 the book circles back to problems and scenarios that were introduced in earlier chapters, showing how the framework of contextual integrity resolves or avoids them. For instance, because contextual integrity demands appropriate flow and not merely control and secrecy, it predicts the behaviors skeptics cite as paradoxical and it also avoids the problem of privacy in public. It readily explains historical and cultural variability, for although the requirement of contextual integrity is universal, variation naturally enters the picture. First, because informational norms are context relative, targeted to specific ends, values, and purposes of these contexts, they must take local requirements of place and time into consideration, at least in the ideal case. Second, relativity is an inherent feature of contexts themselves because different societies evolve different configurations of contexts, resulting in different configurations of actors, attributes, and so on that create the parameters that characterize informational norms. How and why these configurations differ—across distance, time, ethnicity, religion, and nation—is a fascinating question for historians, sociologists, anthropologists, and others, but outside the scope of this book (and this author’s expertise). This all means that historical and cultural variation is not an awkward fact needing explanation but is directly predicted by the framework.

Although chapters 7 and 8 both discuss contextual integrity in terms of real and hypothetical cases, it is Chapter 9 that demonstrates, in detail, the application of contextual integrity to several of the controversial technology-based systems and practices introduced in Part I.

—1  
—0  
—+1

### Scope

I have tried, where possible, to incorporate parallel experiences and significant legal and policy milestones in countries beyond the United States, for example, the European Union Directive, Canadian case law, and the UK experience with CCTV. Readers will see, however, that my reference points on policy and regulation, legal doctrine, and case law are drawn from the U.S. experience. The framework was conceived in the United States and informed by local dramas and rhetoric, public experience, technological milieu, media and landscape, and exemplary or inadequate policy choices, decisions, and practices. Does this mean that contextual integrity is applicable only to the United States? I believe not. It is set forth as a justificatory framework for all people and all societies in which information about people has context-specific function and meaning, and it is governed by norms that systematically reflect these meanings and functions in relation to context-specific ends, purposes, and values. Although many actual examples are drawn from the U.S. experience, there is no reason that key themes and principles should not apply wherever people act and transact in social contexts roughly as I have described.

It must also be acknowledged that the book reflects experience only with technologies of the moment (and the foreseeable but near future). Even in the period over which this book was written, on any given day, month, or year, I could have drawn on different sets of cases, depending on what happened to be front and center at that moment, affected by scientific breakthroughs and historical contingencies. Despite this, I like to think that the framework of contextual integrity transcends particulars of the specific technologies selected for detailed analysis and would apply as well to others. The book's purpose, after all, is to articulate a robust conceptual framework for understanding, evaluating, and resolving critical privacy challenges of the day, past and future.

Finally, the research and scholarship that has most directly influenced this book extends across legal, political, and moral philosophy as well as policy analysis though, to be sure, there are gaps in coverage and, perhaps, disputed interpretations. For readers interested in a broader range of works, the References section provides a useful launch point. Beyond this, there is a growing body of important work on privacy in the empirical social sciences, which deserves more attention than this book has been able to give it. This work is particularly relevant because the framework of contextual integrity asserts

-1—  
o—  
+1—

empirical predictions about actual conditions under which disruptions of flow are likely to draw protest, in contrast with those that are likely to please. Accordingly, important directions for future work on the concept and framework of contextual integrity include checking the plausibility of these predictions against historical findings as well as developing testable hypotheses from them and examining these within research rubrics of the social sciences in natural as well as experimental settings.

—1  
—0  
—+1