Introduction
What's at Stake?

In May 2016, several Danish researchers released data on 70,000 users of the dating website OKCupid. Those of us who have tried online dating know that profiles on OKCupid (or Match, JDate, or eHarmony) are rich in sensitive personal information. The researchers published much of it: usernames, age, gender, and location, as well as sexual orientation, fetishes, religious views, and more. Given the breadth of that information, it wouldn't take much to figure out the identities of those involved. And the researchers neither obtained consent nor anonymized the data.[1]

This data dump posed a privacy puzzle. OKCupid users voluntarily answer hundreds of questions about themselves so the website can use an algorithm to connect them with supposedly compatible matches. The answers are available to all OKCupid members, with some basic information available to nonmembers via a Google search of a person's name. For these reasons, while they may have violated OKCupid's terms of use and ethical canons for research in the social sciences, the researchers felt they were on solid privacy grounds. They did not need to anonymize the data set, they said, because users had provided the information in the first place and other people had already seen it. By any traditional metric, they thought, this information was not private.[2]

Notably, this wasn't an isolated incident. Researchers have mined personal data before.[3] Retailers do it all the time, gathering everything from our browsing histories to Facebook "likes" to target us with advertisements they think we want to see. Google tailors its search results based on what it learns from our behavior across platforms, sometimes discriminating against us in the process.[4] Data brokers amass vast collections of information about us gleaned from across the Web and sell it to the highest bidder. Facebook is steaming ahead with frighteningly accurate facial recognition technology based on the millions of photos we upload for our friends.[5] Retailers analyze our purchasing histories to predict what we will buy next even before we know

it ourselves.[6] And marketers are using our buying patterns and GPS technology to send sale notifications directly to our phones when we pass a brick-and-mortar store.[7]

We have a choice. Our society is at a crossroads. We can live in a world where these activities go on unabated, where stepping outside our homes and using technology are information-sharing events, where the law cannot protect us, and where the only things that are truly private are the things we keep secret. In this world, anyone, whether they are overeager researchers or online advertisers, can use our data because, as a matter of law and social practice, the information is already public. We share our data the moment we sign up for an account, or browse the Internet, or buy a book online.[8] In this world, privacy is dead.

Or, we could live in a world where privacy still matters. In this less totalitarian, more agreeable world, lawyers, judges, policymakers, teachers, students, parents, and technology companies take privacy seriously. Here, privacy has a fighting chance against other values and norms, and information can be shared with others without the entire world looking on.

Today, we seem closer to privacy's death than to its renaissance. Indeed, talking heads have been writing privacy's obituary for years.[9] That is in part because, like the Danish researchers who took a cavalier approach to their subjects' privacy, we have been thinking about privacy too narrowly. Strengthening privacy won't be easy, but getting to a better world starts with changing the way we think about privacy, how we integrate and manage it in our daily lives, and how the law is used to protect it. And that is what this book is about: I want to change the way we think about privacy so we can better leverage law to protect it in a modern world.

We are accustomed to conceptualizing privacy in certain ways. We often think that privacy is about separating from the prying eyes of others or keeping things secret; that's why we draw blinds when we don't want people looking in. Sometimes we associate privacy with certain spaces or property boundaries; what we do "in the privacy of our own homes" or "behind closed doors" or in our bedrooms and bathrooms, for example, is our business. Sometimes we think privacy is bound up with intimacy; that's what makes topics like sex, money, and medical care inherently personal. But I argue that limiting our understanding of privacy to these concepts alone is what allows our data to be mined and used with impunity. These ways of understanding privacy are, at best, incomplete and, at worst, hurtling us toward a dystopian future designed without privacy in mind.

For example, thinking that privacy is synonymous with secrecy could help us when someone hacks our personal network and publishes previously

encrypted files, but it doesn't offer much consolation for users of OKCupid or for a victim of nonconsensual pornography who shared images with a now-ex-boyfriend only to have him upload those pictures to Pornhub. Once we share something, it's no longer secret, and we lose control of it.[10] Similarly, it may be sufficient for a homeowner who notices a drone-mounted camera by her window to think about privacy as bound up with enclosed spaces or property lines, but it would be radically insufficient the moment she stepped outside. And although some of these victims could fall back on the inherent intimacy of the information revealed to explain their feeling that their privacy was invaded, privacy-as-intimacy cannot, on its own, respond to the transfer of personal data among websites, behaviorally targeted advertisements, and "big data" predictive analytics that mine seemingly innocuous and nonintimate information from across the Internet to determine what we see, what we buy, and what we learn.

We need to change our perspective on privacy.

It may sound strange, but privacy is an inherently social concept. The very idea of privacy presumes that we exist in both formal and informal relationships with others: privacy only matters after we share within those relationships. When making sharing decisions, we rely on and develop expectations about what should happen to our information based on the contexts in which we share, thus integrating privacy into our lives relative to other people.[11] As the law professor Robert Post describes, privacy norms "rest[] not upon a perceived opposition between persons and social life, but rather upon their interdependence."[12] Privacy, then, is socially situated. It is not a way to withdraw or to limit our connection to others. It is, at its core, about the social relationships governing disclosure between and among individuals and between users and the platforms that collect, analyze, and manipulate their information for some purpose.[13]

For example, when we share the fact that we are HIV-positive with the 100 members of an HIV support community, we may expect a far greater degree of confidentiality and discretion from them than from two acquaintances at work. When we whisper secrets to a good friend, we expect confidentiality even without a written agreement. We share our bank account numbers with Bank of America's website and expect that it won't be shared with online marketers. And although we may recognize that using the Internet or joining a discount loyalty program requires some disclosure, we share our information with the expectation that it will be used for the specific purpose for which we shared it. What we share, with whom we share it, and how we share it matter. In other words, something about the social context of disclosure is the key to determining what is private and what is not.[14]

That key is trust. Trust is a resource of social capital between or among two or more parties concerning the expectation that others will behave according to accepted norms. It mitigates the vulnerability and power imbalance inherent in disclosure, allowing sharing to occur in the first place. Put another way, disclosures happen in contexts of trust, and trust is what's broken when data collection and use go too far. An information age demands an understanding of information privacy that recognizes that we share a substantial amount of personal data with friends, public and private institutions, websites, and online advertisers. More generally, a doctrine of information privacy must navigate the public/private divide in context, recognizing, among other things, that what we share, when we share, why we share, and with whom we share matter for determining whether disclosed information is still legally protectable as private.[15] And it must not only explain what we consider private, but also why certain things fall into the private sphere and why other things do not.[16] This theory must also reflect how information privacy is implemented on the ground, including how we determine when and what to share, how platforms manipulate us into disclosing more than we might otherwise have wanted, and how, if at all, technology companies embed privacy norms in the data-hungry products they create. And the theory must be administrable, capable of being applied by lawyers and judges in real cases to answer real information privacy questions. Finally, the way we think about privacy has to set us on a better path, one that not only helps privacy thrive in a modern world, but also has positive effects on society as a whole.

Because we share when we trust, I argue that we should start talking about, thinking through, and operationalizing information privacy as a social norm based on trust. In the context of information sharing, trust gives us the ability to live with and minimize the vulnerability inherent in sharing by relying on expectations of confidentiality and discretion. Indeed, all disclosures create vulnerability and imbalances of power. Elsewhere, as in doctor-patient or attorney-client relationships, where significant disclosures create similar power imbalances, we manage those risks with strong trust norms and powerful legal tools that protect and repair disclosure relationships. Reinvigorating information privacy requires similar norms and legal weapons, as well. So, when we share information with others in contexts of trust, that information should be protected as private. I call this argument *privacy-as-trust*, and, like other trust doctrines in the law, it allows disclosure to occur in safe environments buttressed by concurrent norms of confidentiality and discretion.

By the end of this book, my hope is that we will start considering trust as an important part of our notion of information privacy. More specifically, my goal is to argue that we should conceptualize information privacy in terms of

relationships of trust and leverage law to protect those relationships. This, however, does not mean that all other visions of privacy are useless. Important rights-based concepts of privacy are not wrong; they are just incomplete. On their own, they have difficulty answering some modern privacy questions posed by new technologies like predictive analytics, social robots, and ongoing and pervasive data collection. Privacy-as-trust can help get us on the path to a better world where privacy not only exists, but thrives, and where society benefits from a rejuvenation and strengthening of trust norms among individuals and between individuals and data collectors.

It is also important to note what this book is not about. Privacy takes on many forms, in different contexts, with a variety of bogeymen ready to break or erode it. This project is about *information* privacy, generally, and privacy in times of disclosure, specifically. It is primarily about the ways in which we interact and share information with, and are vulnerable to, other private actors – other individuals and technology companies, for example – rather than government agents. That is not to say that conceptualizing privacy as based on relationships of trust is necessarily silent or unhelpful in a variety of contexts. But those extrapolations and extensions are for another book.

I construct my argument in three stages. Part I is about where we have been; it develops and then critiques the many theories of privacy that dominate current privacy scholarship, showing how each of them is a variant on the same theme and has helped bring us to where we are today. Part II is about the theory of privacy-as-trust itself; it teases out the definition of trust, provides empirical evidence in support of the relationship between trust and disclosure, and shows how privacy-as-trust is already being operationalized on the ground. It argues that trust must be part of our understanding of privacy, as a result. Part III is about the better world with trust in it. I apply privacy-as-trust to several vexing questions of privacy and information law, and show the contrast between conventional and trust-based approaches. In all cases, understanding privacy as bound up with the concept of trust brings about a better, more just world where privacy is a strong social value.

PART I: WHAT DO WE MEAN BY "PRIVACY"?

For many, privacy is about choice, autonomy, and individual freedom. It encompasses the individual's right to determine what she will keep hidden and what, how, when, and to whom she will disclose personal information. Privacy is her respite from the prying, conformist eyes of the rest of the world and her expectation that things about herself she wants to keep private will remain so. I will call these ideas the *rights conceptions of privacy* to evoke their

Lockean and Kantian foundations. And they can be divided into two categories. In Chapter 1, I discuss the definitions of privacy that are based on negative rights, or those that see the private sphere as a place of freedom *from* something. These notions of privacy include elements of seclusion and private spaces, as well as conceptions based on the sanctity of private things, like discrediting secrets or intimate information. Common to these ways of thinking about privacy is an element of separation, suggesting that they provide freedom from the public eye. Chapter 2 discusses the second category of rights-based definitions of privacy. These conceptualizations retain the assumption of separation, but use it for a different purpose – namely, for the opportunity to grow, develop, and realize our full potential as free persons. It conceives of privacy as affirmatively *for* the full realization of the liberal, autonomous self.[17]

The rights conceptions of privacy pervade privacy rhetoric, scholarship, and judicial decisions. They are the dominant ways we approach privacy problems today. They are, however, incomplete. They miss the fact that information privacy norms are triggered by disclosure. And disclosure is an essentially social behavior: once we share, we trade control of our information for reliance on powerful social norms, or background social rules that feed into our expectations of what should happen with our personal data. Privacy centered solely on the individual ignores those social norms even though they are not only essential to sharing but have positive effects on social solidarity. Without them, we risk narrowing privacy into oblivion.

Like the work of Robert Merton,[18] Michel Foucault,[19] Helen Nissenbaum,[20] and others, privacy-as-trust approaches information privacy and disclosure from a social perspective. Privacy-as-trust recognizes that information privacy is not about excluding others, but rather about regulating the flow of information to some, restricting it from some, and opening it up to others. This essential understanding about privacy's social role is not new, and Chapter 3 focuses on describing the development of social theories of privacy over the last 50 years. In that chapter, I argue that social theories of privacy to date may have recognized that privacy is what manages information flow in context, but they inadequately respond to the power dynamics at play in disclosure. That is the role of trust.

PART II: TRUST AND PRIVACY

Disclosure and privacy govern our relationships with others (persons as well as technology platforms); as such, they are social phenomena. Trust is the link between them. And strong trust norms are what allow sharing and social interaction to occur.[21]

Particular social trust is the "favourable expectation regarding other people's actions and intentions," or the belief that others will behave in a predictable manner.[22] It "begins where knowledge ends"[23] and is the mutual "faithfulness" on which all social interaction depends.[24] For example, when an individual speaks with relative strangers in a support group like Alcoholics Anonymous, she trusts that they will not divulge her secrets. Trust, therefore, includes a willingness to accept some risk and vulnerability toward others and steps in to grease the wheels of social activity.[25] I cannot know for certain that my fellow support group members will keep my confidences, so trust allows me to interact with them, disclose information, and rely on discretion with confidence. And I earn all sorts of positive rewards as a result.[26]

It makes sense, then, to turn to trust when thinking about what motivates us to share personal information online and what governs the related privacy norms in social interaction: Alice shares information with Brady because Alice trusts Brady with that information; the applicable norms – confidentiality and discretion – give Alice the confidence and comfort to share with Brady, mitigating the vulnerability inherent in someone else having access to her home. The same mechanism is at play when we share information with lawyers, doctors, and financial planners: strong trust norms, backed by tradition, professional standards, and the law, give us the confidence and comfort to share. Despite the intuitive appeal of that mechanism, particular social trust has been, at best, a silent undercurrent in a growing literature on our propensity to disclose personal information. Part II of this book teases out this privacy, sharing, and trust relationship.

The theory of privacy-as-trust is the subject of Chapters 4 and 5. Privacy-as-trust posits that information disclosed in contexts defined by trust should be legally protected as private. It is not an attempt at a unitary, a priori definition of privacy that applies to all situations.[27] But privacy-as-trust does give us a way of understanding how private disclosure contexts vary from context to context and why certain uses of data strike us as invasive and unfair. Thinking about information privacy as based on relationships of trust means several things. It means seeing privacy as something that can foster disclosure by mitigating the risks inherent in disclosure and rebalancing power between sharers and audiences. It means looking at the context of disclosure to determine the difference between public and private information. It means considering both norms at the time of disclosure and any background that has an impact on future expectations. And it means asking how the law can be used to strengthen relationships of trust between parties and equalizing the power imbalances that come with sharing. Doing this will have significant value to society.

Chapter 4 also offers some evidence that privacy-as-trust reflects how we operationalize privacy and disclosure decisions in practice. An empirical study of Facebook users, summarized here, suggests that trust is a key factor in users' decisions to share personal information on the platform. And, as scholars have shown, many companies with strong privacy leaders at the top think about their privacy obligations as protecting and fostering trust between the company and its customers.[28] Therefore, if trust is defining our understanding of privacy on the ground, perhaps the law and privacy theory can catch up.

PART III: PRIVACY-AS-TRUST IN ACTION

The balance of the book considers what information privacy law would look like if we applied privacy-as-trust to several ongoing privacy and information law controversies. There are five chapters in this section, each of which uses case studies to show how a privacy law regime based on trust would look different than the status quo. In each case, some amount of disclosure causes risk, vulnerability, and a loss of power; privacy-as-trust restores the trust norms that protect those disclosures in the first place.

Chapter 6 starts at the macro level, considering Internet platforms' obligations and responsibilities. The current regime, which requires data collectors do little more than provide us with notice of what information data they collect and what they do with it after collection, is based on the idea, discussed in Chapter 2, that privacy is about the freedom to choose when and how to disclose personal information. As many scholars have argued, however, this "notice-and-choice" approach is hopelessly flawed and inadequate. It gives users little to no help when making disclosure decisions, and it offers even less protection when Internet companies use our data in unexpected and invasive ways. This is especially problematic where web platforms use artificial intelligence (AI) or complex algorithms to learn about us, predict the things we want to see, and mediate our online experiences. A reorientation of privacy law around principles of trust would address these gaps, providing the necessary theoretical justification for holding data collectors to certain fiduciary responsibilities of loyalty. As the legal scholars Jack Balkin, Jonathan Zittrain, and others have argued, this would protect us from Internet platforms that are already inducing our trust, taking our data, and harming us for their own profit.

Chapter 7 goes from macro to micro, applying privacy-as-trust to several cases about the wide dissemination of information previously disclosed under limited circumstances. These cases apply the current privacy torts, including intrusion upon seclusion and public disclosure of private facts, which require

judges to determine the difference between public and private information. But when judges, as many do today, define privacy as synonymous with secrecy, victims of privacy invasions are left out in the cold. Privacy-as-trust, like other social theories of privacy, recognizes that privacy exists post-disclosure and provides judges with clear, easy-to-apply questions to make more nuanced decisions.

Chapter 8 argues that privacy-as-trust forces us to think differently about privacy harms. Currently, most scholars and judges see invasions of privacy as attacks on the individual. Privacy-as-trust recognizes that because trust is what facilitates and regulates information flows, injuries to information privacy are injuries to the norms of social interaction. This opens up a new avenue for protecting personal privacy: a robust tort of breach of confidentiality. Traditionally marginalized in American law, the tort is perfectly suited to protecting the privacy of previously disclosed information. This chapter looks at one type of cyberharassment – nonconsensual pornography or so-called revenge porn – as an illustrative case study.

Chapter 9 steps outside the confines of privacy to show that privacy-as-trust can be used as a more general theory of information flows. Using a case study of patent law's public use bar, which prevents inventors from securing a patent if they have shared their invention with the public more than one year prior to application, this chapter makes several arguments. Judges today apply some of the same rights-based principles discussed in Chapters 1 and 2 to define "public" in the public use bar. This has the perverse effect of privileging wealthy corporate inventors because it ignores the unique social determinants of information flows among solo entrepreneurs. Privacy-as-trust not only addresses that imbalance but also provides a clear, administrable, and fair way to distinguish public from private inventions.

Finally, Chapter 10 discusses social robots. Social robots – like Sony's Aibo dog or Kaspar, a machine with humanlike qualities designed by the University of Hertfordshire to help children with autism learn how to respond to others – are machines that interact with humans on a social level. They pose special legal challenges that traditional understandings of privacy cannot comprehend. Social robots are both wonderful and insidious, and their dangers are directly related to their benefits: while helping us meet foundational human needs of companionship, friendship, and emotional connectedness, they distract us as they sweep in troves of personal data. Plus, as machines with human*ish* tendencies, they elicit more emotional responses than rational ones. This makes us vulnerable, especially in a conventional privacy world where mere use of a technology product is considered consent to ongoing data collection. Privacy-as-trust explains the dangers of social robots – we are

primed to trust them and, as result, eager to share – and suggests ways to protect ourselves in the process.

I conclude by summarizing my argument and suggesting avenues for future research. Privacy-as-trust is not about keeping more information private or making more information public. Thinking about privacy as a social norm based on trust can help individuals protect themselves against invasions of privacy. It can also foster disclosure where needed. It fosters productive relationships, both commercial and social. And it fosters powerful trust norms that could bring us closer to each other and to technology companies who act responsibly with our data. If we know that the websites we use respect our trust-based disclosure expectations, we may feel comfortable sharing more information to enhance our online experiences. We just need a new way of thinking about privacy that works for us, not just for data collectors. This book offers that opportunity.