

The Economics of Privacy

Introduction

By Avi Goldfarb and Catherine Tucker

Digital economics focuses on how the economy is transformed when information can be stored in bits and not atoms. This advance has reduced the cost of collecting, storing, and processing data. Firms and governments can observe individual-level behavior at a detailed level, and use this information to provide better products and services; however, the ability to access this information raises privacy concerns.

This volume provides a summary of the research to date on the economics of privacy, and identifies a number of open questions. The five chapters in this volume are based on presentations that the authors gave to the NBER PhD tutorial on the Economics of Privacy, held in October 2022. They provide different perspectives on the role of economists and economics research in understanding the benefits of privacy and the benefits of data flows.

Several common themes emerge across the chapters.

One common theme is the challenge in applying the Coase theorem to privacy. To many economists, it is compelling to assume that privacy concerns can be solved by the market because the parties should be able to negotiate a price for exchange of data if property rights are well-defined and transaction costs are low. Each chapter describes situations where the Coase theorem breaks down.

A second common theme relates to measurement. It is relatively easy to measure the costs of privacy in terms of innovation or competition and so much of the existing literature has documented such costs. It is more difficult to measure the benefits of privacy, particularly when those benefits are about intrinsic preferences rather than a direct and measurable market harm.

A third common theme is the difficulty in defining privacy through an economics lens. The broader academic and legal literature includes many distinct definitions of privacy. With our emphasis on formal models, economists must select a definition and formalize it. This creates challenges in the literature as people use the same word, “privacy”, to mean different things.

We next briefly describe each chapter in turn.

In Chapter 1, Catherine Tucker focuses on identifying the open research questions, linking the economics perspective to changes in other fields. The chapter defines privacy using the definition from Warren and Brandeis (1890) as “The right to be let alone”. The chapter comprehensively lists a set of outstanding questions, and discusses the role of the economist’s perspective for addressing these questions. Key open questions include the value of privacy, the role of property rights and data markets, the relationship between privacy and inequality, and the political economy of privacy.

Alessandro Acquisti takes a different approach in Chapter 2, arguing that the typical economic approach is flawed in the context of privacy. While the chapter's core point matches a point made by Tucker, that we need more research on the benefits of privacy, Acquisti argues that the economic approach is limited because it focuses on what we can measure and what markets can solve. The chapter defines privacy using Altman's (1976) definition as "a dynamic and dialective process through which individuals contextually manage the boundaries between the self and others." Ultimately, the chapter concludes that economists would benefit more from reading the privacy literature outside of economics, and provides a brief summary of key ideas in that literature.

In Chapter 3, Alessandro Bonatti provides a formal approach to understanding how privacy, defined as a restriction on data flows, affects market outcomes. The chapter begins by providing a comprehensive theoretical framework for understanding privacy. Then the chapter explains how the broader economics literature on privacy can be interpreted using this model. A key idea emerging in this literature, and central to the chapter, is the role of data externalities. The provision of data by one person can affect others. One aspect of this is the social dimension of data, as data about other people can be informative about one's own behavior. Like Tucker in Chapter 1, this chapter identifies open questions related to markets for privacy and to political economy.

Garrett Johnson's chapter provides a detailed analysis of a particular regulation: Europe's General Data Protection Regulation (or GDPR). The chapter provides a detailed description of what is likely the most important digital privacy regulation to date. In the process, it highlights several ongoing challenges in empirical privacy research. Regulation has benefits and costs, and that costs are more straightforward to measure in terms of reduced competition and reduced innovation. The chapter highlights how the regulation as written can be different from the regulation as implemented. This distinction between the law and its enforcement can lead to a misinterpretation of the effects of privacy regulation on behavior and on market outcomes. Overall, the chapter provides a useful framework for economists looking to examine how changes in privacy policies by firms or governments might affect economic outcomes.

Fifth, and finally, Amalia Miller describes the role of privacy in healthcare. She notes that health is a particularly interesting area to study privacy because there are extraordinary benefits of data flows and an extraordinary potential for harm. The benefits are clear as improved healthcare means longer and healthier lives. Better information about individual patients could lead to improved care. Better aggregated health information could lead to better diagnostics and treatments. The chapter also lists a variety of harms. Direct harms include a feeling of shame or embarrassment, a feeling of betrayal, and a feeling of invasion or loss of freedom. Indirect market harms include labor market harms, insurance, personalized pricing (if higher), and targeted advertising (if manipulative, annoying, or intrusive). Indirect non-market harms include social stigma, harms to reputation, increased risk of identity theft or fraud, and increased legal risk. While the chapter emphasizes these harms in the context of health information, the division into direct harms, indirect market harms, and indirect non-market harms provides a useful framework for the broader economics of privacy literature. This perspective is grounded in economic models of the costs and benefits of data flows, and the potential for externalities to undermine standard economic approaches. The chapter also notes that there is a great deal of health privacy

regulation in place in the United States and around the world, making health a fruitful area for empirical work. The last line of chapter 5 summarizes the challenges for decision-making around privacy, highlighting the “fundamental tradeoffs between preserving privacy and harnessing the value of IT and data-driven innovation”.

This volume aims to summarize the key open questions in the economics of privacy as of 2023. In doing so, the chapters provide different frameworks that economists can use to undertake research on privacy.