

The Economics of Privacy at a Crossroads

Alessandro Acquisti¹

Carnegie Mellon University

Draft, July 2023

NBER Tutorial on the Economics of Privacy

By several accounts, the economics of privacy has grown into a remarkably successful field of research. As the means of collecting and using individuals' data have expanded, so has the body of work investigating trade-offs associated with those data flows. The number of scholars working in the area has grown, much like the breadth of topics investigated. References to the economic value of personal data have become common in policy and regulation, and so have mentions of economic dimensions of privacy problems.

Thinly veiled underneath those successes, however, lies a less encouraging trend. In this manuscript, I argue that the very success of the economics of privacy has laid the foundation for a potentially adverse effect on the public debate around privacy. Economic arguments have become central to the debate around privacy. When used as complements to considerations less amenable to economic quantification, those arguments are valuable tools: they capture a portion of the multiform implications of evolving privacy boundaries. When, instead, economic arguments crowd out those other non-economic considerations from the public discourse around privacy, problematic scenarios arise. In one scenario, the economic analysis of privacy will keep growing in influence, but its overly narrow conception of privacy will impoverish rather than augment the depth of the debate around privacy. In a second scenario, less likely but equally problematic, the economics of privacy will progressively undermine its own relevance by failing to account for the complexity and nuance of modern privacy problems.

There is a third scenario—one this manuscript explores. The economics of privacy may expand its horizons and relevance both by considering economic dimensions and research questions that have so far received limited attention, and by accounting for the broader scholarship on privacy coming from other disciplines. As a complement to the contributions of other fields, rather than a substitute for them, the economics of privacy may keep thriving and remain a useful tool for debate and policymaking.

My argument, and this manuscript, proceed in three steps roughly focusing on the past, present, and possible future of the economics of privacy. In Section 1 I focus on the past. I review the rise of this field of research up to the current days and celebrate its successes. In Section 2 I

¹ I gratefully acknowledge support from the MacArthur Foundation through grant 22-2203-156318-TPI. I am also grateful to several scholars for comments and critiques, including Idris Adjerid, Laura Brandimarte, Cristobal Cheyre, James Cooper, Brett Frischmann, Avi Goldfarb, Chris Hoofnagle, Brian Kovak, Michael Kummer, Tesary Lin, Jonathan Mayer, Klaus Miller, Verina Que, Ananya Sen, Priya Shah, Daniel Solove, Ryan Steed, Andrew Stivers, Lior Strahilevitz, Catherine Tucker, Joel Waldfogel, Frederik Zuideveer Borgesius, one reviewer, and participants in workshops and seminars at HEC Paris, Massachusetts Institute of Technology, New York University, Princeton University, University of Maryland, University of Minnesota, and the ZEW Conference on the Economics of Information and Communication Technologies.

take stock of the present, and focus on the unintended consequences of those successes. I consider the shortcomings of the economics of privacy arising from its misconstruction or dismissal of critical privacy theories from other social sciences. In Section 3, I consider a possible alternative future for the economics of privacy. I propose ways of framing the economic debate around privacy that deviate from the focus of much (albeit not all) current research, which tends to concentrate on the economic costs of privacy protection at the expense of a richer array of yet unanswered questions.

Section 1: The Rise of the Economics of Privacy

The economics of privacy is not a novel field of research. It boasts a venerable pedigree.² A wave of economic analyses of privacy started appearing near the end of the 1970s and the start of the 1980s. Chicago scholars interested in economics and law, such as Posner and Stigler, produced several of those analyses. They were the intellectual “pioneers”—in Hirshleifer’s (1980) wording—who had discovered a “new territory [...] the intellectual continent we call ‘privacy’” (p. 649). It is not ungenerous to describe (using Hirshleifer 1980’s words) those pioneers’ views as “hostile” (p. 650) towards privacy. Posner (1977, 1978, 1981) identified privacy, from an economic perspective, as the concealment of information (in particular, *negative* information), and surmised that regulations intended to protect privacy would ultimately be redistributive and result in economic inefficiencies. Stigler (1980) believed that privacy “connotes the *restriction* of the collection or use of information about a person or corporation” (p. 625; emphasis added). He found the spur of new interest in it “paradoxical, for the average citizen has more privacy—more areas of his life in which his behavior is not known by his fellows—than ever before” (p. 623). Not everyone agreed with those views. Hirshleifer (1980) countered that privacy was more than a restriction on data collection; it was about *autonomy within society*. He wondered whether the new continent that economics had discovered was not, in fact, merely a *peninsula* that those economic pioneers had mistaken for the mainland. In some sense, the dispute Posner and Hirshleifer commenced over four decades ago has never been truly resolved. Its relevance to the current debate around privacy will become apparent as this manuscript progresses.

After its first wave of research output, the economics of privacy went largely dormant until the mid-1990s, when a new generation of economists such as Varian, Noam, and Laudén started rediscovering the topic. The reasons why economic interest in privacy reemerged at that time seem clear in retrospect. The information technology revolution was transforming (digitizing) the collection and use of personal data, and the World Wide Web was developing. These scholars captured the impending economic implications of those changes. Varian (1996) diagnosed the link between economics and technology at the root of the modern privacy problem: data that was already theoretically public (or at least accessible) in physical format becomes much cheaper to capture, store, access, and share once digitized, and thus “more” public; as its price lowers, quantity demanded increases. Noam (1995) wrote about the economic interpretation of

² Others have reviewed its evolution in detail (among them Hui and Png 2006; Acquisti, Taylor, and Wagman 2016; Cecere, Le Guel, Manant, and Soulié 2017; Goldfarb and Que 2023). Here, I only highlight a few key milestones.

encryption and data protection. And Laudon (1996) was arguably the first economist to lay out the idea of data markets through which individuals could one day trade rights over their personal data. That idea has taken different manifestations in the roughly twenty-five or so years since it was first proposed, and has been the subject of numerous proposals (from data dividends to data as labor; see Arrieta-Ibarra et al. 2018).³

The scholars who contributed to the economics of privacy in the mid-1990s added nuances to the minimalist view of privacy espoused by Chicago School scholars in the 1980s. For instance, Varian (1996) noted that individuals may strategically prefer to share some personal information while protecting other, not necessarily negative, information: the same consumer may want her preferences to be shared with a merchant (to get personalized offers), but not her reservation price (to avoid first-degree price discrimination).

Boosted by tectonic changes brought about by the development of the Internet, the field of the economics of privacy eventually took off. Over the last two decades, the costs of data collection, storage, and computation kept falling; while the sophistication of statistical techniques for inferential data analysis kept rising. These combined trends led to the development of strategies for data monetization and to the identification of personal data as an economic asset. This, in turn, spurred novel products and services, which created more data, which generated more value, which attracted more investments, and so forth. As this feedback cycle developed, the data economy grew, and so did the economics of privacy. Economically-informed position papers from the 1990s were replaced by analytical models; empirical studies started testing the theories; field and lab experiments became commonplace; and the specific topics of investigation under the vast umbrella of “privacy” economic research started expanding and diversifying—although they remained mostly tied to the informational dimension of privacy. While in the early 2000s, much research on this topic focused on data breaches and price discrimination, the topics covered have multiplied over time—from the relationships between data and competition and antitrust to the creation of data markets; from the link between privacy regulation and innovation to data-driven algorithmic bias; from experiments on consumer data valuation to studies of behavioral factors affecting privacy decision-making. The number of scholars authoring manuscripts in this field has increased, and their backgrounds have become more diverse—from mainstream economics to marketing; from information systems to computer science. The number of unique outlets (conferences and journals) publishing work in the field has also increased; it has become more common to see articles published in traditional premiere economic outlets such as the *American Economic Review*, the *Journal of Political*

³ In our 2016 review of the economics of privacy in the *Journal of Economic Literature* (Acquisti et al. 2016), as well as in other recent pieces (Spiekermann, Acquisti, Böhme, and Hui 2015; Acquisti, Brandimarte, and Loewenstein 2020), we discuss some of the reasons why, although consumer data is now an asset explicitly or implicitly traded in a myriad of ways, personal data markets such as those envisioned by Laudon remain elusive—notwithstanding widespread scholarly and commercial interest. Central among those reasons are the absence of regulation creating well-defined property rights over personal information, as well as the fact that most of the more valuable personal data is not static (e.g., a person’s gender) but dynamically co-created by the data subject and platforms or services the subject interacts with (e.g., a person’s preferences, as revealed by her most recent search query or visited website). Absent regulation explicitly giving individuals rights over their personal data (including co-created data), platforms maintain economic control over it, undermining consumers’ ability to leverage parallel “data markets” to protect (or merely commercially benefit from trades over) their personal information.

Economy, and the *RAND Journal of Economics*. There are anecdotal yet meaningful signals of relevance, too: the publication in the *Journal of Economic Literature* of a review of the field, and the hosting at the *National Bureau of Economic Research* of a Workshop (in May 2022) and a Tutorial (in November 2022) on the economics of privacy—both organized by two scholars who have been at the forefront of the revival of the field, Professors Goldfarb and Tucker. In a nutshell, one could say that over the course of four decades, the economics of privacy has attained a meaningful role within the economics mainstream.

Section 2: Where is the Economics of Privacy Going?

As economists, we like to talk about unintended consequences, often referring to undesired ramifications of regulatory interventions in the market. In this section, I discuss the unintended consequences of the success the economics of privacy has experienced as a field of research. I start by comparing the way economists have traditionally construed privacy to the conception of privacy developed by influential privacy scholars who have worked outside of the realm of economics (Section 2.1). I argue that as economists we have, by and large, adopted a reductionist view of privacy that overlooks the richness and nuance of the contemporary debate around privacy. Next (Section 2.2), I discuss the unintended consequences of that approach. They include an outsized focus on estimating the costs of privacy regulation at the expense of a more comprehensive analysis of the diverse trade-offs of privacy; a lack of attention to the many consumer harms of privacy intrusions; and a misapprehension of the lessons of behavioral privacy research. Ultimately, the rigorous but narrow approach to privacy (pioneered by Posner and Stigler but still - I argue - influential in today's economic research) carries the risk that economic arguments may crowd out of the public debate the discussion of privacy dimensions that are not grounded on economic analysis and yet are no less important.

2.1. How Privacy Economists Think of Privacy and How Privacy Scholars Outside Economics Think of Privacy

The essays Posner wrote concerning privacy in 1977, 1978, and 1981 proved over time remarkably influential, and not merely because of their citation count. Notwithstanding the dramatic growth and evolution of this field of research, the influence of those essays can still be detected, today, in the framing and scope of much (but not all) economic research in the field.

Posner makes four remarkable points in the first pages of his 1981 article. First, after having acknowledged different interpretations of the term “privacy,” Posner identifies one as the most deserving of economic attention: concealment. Second, Posner narrows down the scope of fruitful economic analysis of privacy to the study of *concealment of information*—and, more specifically, *negative information*: to the extent that an individual is deficient in some characteristics (in Posner's example, an employee may be deficient in terms of diligence, loyalty, or mental health), she will have an incentive “to conceal these deficiencies” (p. 405) and to “invoke a ‘right of privacy.’” If privacy is the concealment (first point) of information, and specifically of negative information (second point), the third point logically follows: “[b]y reducing

the amount of information available to the ‘buyer’” (in Posner’s labor market example, an employer), privacy “reduces the efficiency of that market” (p. 405). Posner’s fourth point concerns consumer privacy behavior, which he deems consistent with theories of rational choice: “the literature on the economics of nonmarket behavior suggests that people are rational even in nonmarket transactions [...] [t]herefore, there seems to be no solid basis for questioning the competence of individuals to attach appropriate (which will often be slight) weight to private information—at least if ‘appropriate’ is equated with ‘efficient’” (p. 406).

I am not going to claim that Posner’s construction of privacy as the concealment of *negative* information (his second point) still influences today’s economic scholarship around privacy.⁴ The notion of privacy as being about something to hide has been repeatedly debunked (Solove 2007), and nowadays most economists, I venture, would reject the reasoning behind that claim. I am interested, instead, in discussing how Posner’s other claims have influenced economic research in this area (including, for full disclosure, my own), and how they compare to the theories and findings of some prominent privacy scholars outside of the economic domain. I am also not going to claim that the entirety of the economic discipline still endorses in lockstep all of Posner’s other three points. Below, I attempt to highlight both cases where the field evolved and diversified and cases in which Posner’s conception of privacy still profoundly permeates our writings—including the research questions we tackle and the implications we draw.

A first difference between the mainstream economic approach to privacy and that of privacy scholars in other social sciences pertains to the very definition of privacy. Posner proposed that *concealment* (of information) was privacy’s most interesting meaning from an economic standpoint. Since then, explicit or implied references to privacy as concealment, restriction, or protection of information have remained common in both analytical economic perspectives on privacy and in empirical economic works. As Lin (2022, p. 665) recently notes: “economists often think of privacy preference as generated from the need to *protect one’s private information* in market exchanges” (emphasis added). To be clear, in some cases the field has evolved from the Posnerian view. For instance, Noam (1995) referred to privacy as “an interaction, in which the rights of different parties collide” (p. 51). And Jin and Stivers (2017) drew the key distinction between privacy processes and privacy outcomes, noting that consumers “want [...] to have a certain amount of *control* over the flow [of individual information]” (emphasis added). But the focus on information restriction still trickles up in our economic writings. For instance, Jin and Stivers also define “an individual’s privacy outcome” as “the realized *restriction* on the flow and use of information” (p.1, emphasis added); “[a]n entity has more privacy as the flow and use of information about it is more restricted” (p. 5).⁵ And in a remarkable study of the use of electronic medical records to prevent AIDS deaths by enabling patient tracing, presented at the

⁴ Although it still sporadically shows up in the public debate around privacy. In a 2013 interview, Eric Schmidt (then Google’s CEO) famously answered a question concerning Google’s privacy controversies by stating: “If you have something that you don’t want anyone to know, maybe you shouldn’t be doing it in the first place.” See <https://www.cnbc.com/inside-the-mind-of-google/>.

⁵ Dr. Stivers however notes, in reference to his 2017 manuscript: “Privacy, in our view, was a particular kind of outcome that we were in part trying to point out was too limited [...] Posner defines privacy as concealment, and throws out the rest, but we make the distinction precisely to point out that concealment isn’t the issue, its control and the opening and closing, as appropriate” (personal communication with the author).

first NBER Workshop on the Economics of Privacy, Derksen, McGahan, and Pongeluppe (2022) construe privacy in terms of patients' refusal to be traced for medical purposes (hence the title: "Privacy at What Cost? Saving the Lives of HIV Patients With Electronic Medical Records").

For most privacy scholars (by which I refer to scholars who predominantly study privacy across disciplines as diverse as sociology, psychology, behavioral research, communication, or philosophy), privacy may *include* concealment or protection as one means, but both its means and its ends are broader, more nuanced, and ultimately different from concealment. Across other social sciences, privacy is not just about concealment or exclusion. Privacy has been linked to (and defined in terms of) control, boundary regulation, and more (see Altman 1976 for a comparison of privacy theories; see taxonomies and references in Solove 2006 and Acquisti, Brandimarte, and Loewenstein 2015). Even when narrowly applied to information, "control" is construed as more than protection; it implies the ability both to protect and to share about oneself (Westin 1967). In essence, within much social science research, privacy is not a static condition of hiding but rather—as the American social psychologist Irwin Altman (1976) put it—a process of *boundary regulation*. Under this perspective, privacy is a dynamic and dialectic process through which individuals contextually manage the boundaries between the self and others. It is dynamic because the process changes and evolves according to context. It is dialectic because both sharing and protecting (for instance, personal information) can be privacy management behaviors. When a person chooses to share a secret with a friend to get her advice, that person is engaging in boundary regulation, as they selectively opted to share this information only with her. If the friend later betrays the person's trust (for instance, she gossips about that secret), that is the moment the boundary has been broken and the person's privacy violated.

The difference between concealment and control (or regulation) may appear too abstract and ambiguous. As economists, we may feel queasy about studying concepts seemingly as intangible as the "regulation of boundaries," and uncomfortable with the multitude of dimensions of privacy enumerated across the privacy literature. That multidimensionality may appear to lack the precision and rigor we need for analytical research. And yet the distinction is highly consequential in terms of how consumer behavior around privacy is (mis)interpreted within economics. Through the lens of privacy scholarship, for instance, HIV patients being alarmed about medical tracing and rejecting electronic medical records (as in the findings by Derksen et al. 2022 cited above) is not a failure of too much privacy but too little: when patients cannot trust how their data will be used, they avert sharing; if they could trust that their data would be protected and only used for the intended medical treatments, they would be more likely to share it with doctors and benefit from doing so.⁶ By missing those nuances, as economists we risk self-selecting into an overly constrained analysis of the phenomenon we purport to study, or, worse, we risk ascribing to privacy merits or faults that may not be its own.⁷

⁶ In a study of state genetic privacy laws, Miller and Tucker (2018) find that approaches that give users control over redisclosure *encourage* individuals to obtain genetic testing, while notifications deter them.

⁷ With some exceptions (for instance, Laudon 1996 citing Westin 1967, or Bleier, Goldfarb, and Tucker 2020 citing Nissenbaum 2004), it is telling—and alarming—that references to the writings of some of the most influential scholars and theorists of privacy, such as Westin, Altman, Petronio, or Nissenbaum—to mention a few—are rare in economic

A second notable difference between economic and privacy scholarship follows as a corollary of the definitional difference I have just highlighted. It pertains to the *scope* of the investigable privacy domain. Consistent with Posner's focus on privacy as concealment of information, most of the economic scholarship has concentrated on the study of *protection of data*. While the specific application areas have expanded through the years (to include medical privacy, technological innovation, algorithmic bias, online advertising, and so forth), the modeling literature has tended to focus on the collection of consumer preferences, traits, or reservation prices across various application domains and, with some exceptions (see Section 2.2.1), the empirical literature has focused on the economic ramifications of curtailing access to those data through regulation, self-regulation, or technology.

Privacy economic research has good reasons to focus on information and data. Information assets have become central to the economic calculus of people and organizations, and the novel privacy concerns that have arisen in recent decades are, at least on first analysis, informational concerns. However, in doing so, privacy economic scholarship deviates from the rest of social science privacy research and (I argue in Section 2.2.2) misses the bulk of harm individuals and society can suffer when privacy is mismanaged. Privacy scholars do not identify privacy with data, and Altman's theory of boundary regulation does not merely apply to informational boundaries. Multiple boundaries exist between the self and others, including spatial, bodily, and decisional. Those boundaries can take different embodiments depending on context; what they have in common is the alternating of the opening and closing of the self to others. This is why privacy, for privacy scholars, is tied to—and sometimes a necessary antecedent for—other concepts such as freedom (including bodily freedom), dignity, liberty, autonomy (including decisional autonomy), and so forth. These other dimensions of privacy are hard to quantify and, if not entirely ignored, are thus to a great extent sidestepped in the economic debate around privacy.⁸ And yet, sidestepping that definitional richness, I argue in Section 2.2, is why economic analysis fails to fully grasp the role and impact of privacy in society.

A third difference pertains to the divide between some economists' interpretation of consumer privacy behavior (and decision-making) and that of behaviorally-focused privacy scholars, and their differing estimations of consumer demand for privacy. Posner and Stigler looked at consumers' disclosure decisions as economically rational processes, where individuals strategically signal positive traits but hide negative ones. The belief that consumers can make *economically* rational privacy decisions is still reflected today in the interest some economists have demonstrated toward data markets or toward privacy policymaking that favors informational interventions to assist consumers in navigating privacy trade-offs in the market. Even economists who acknowledge the challenge raised by informational asymmetries (for instance, Jin and Stivers 2017) highlight the role of informational interventions in ameliorating

writings. I too came to appreciate the significance of Altman's writings for economic research only following the completion of my doctoral studies.

⁸ Prudent and perhaps intentional as such sidestepping may sound from an economist's perspective, it raises one of key issues this manuscript attempts to tackle: when we use economics to study privacy, are we aware that we may be missing the forest for the trees?

consumer privacy choice in the marketplace. That belief is also reflected in empirical research that attempts to demonstrate the stability of privacy preferences and the economic rationality of privacy decision-making (Lee and Weber 2021). That belief, in turn, informs how the results of empirical consumer research are interpreted in terms of consumer demand for privacy. As economists, we are trained in the concept of *revealed preferences*. If privacy is narrowly construed as *protection* of personal information, and if privacy behavior is (assumed to be) economically rational, then a revealed preferences perspective would lead us to interpret the abundant evidence of widespread public disclosures (facilitated by social media and embraced by a significant portion of the world population) as realizations of market equilibria that reflect consumers' "true" underlying preferences for privacy. In turn, such evidence could then be interpreted as proof that individuals do not care for privacy that much, and that (regulatory) interventions in this domain are therefore not advisable or required. Results from experiments where participants willingly departed with their personal information in exchange for tiny rewards (Athey, Catalini, and Tucker 2017; Grossklags and Acquisti 2007) may be interpreted as ultimately supporting these conclusions.

In contrast, behavioral privacy research presents evidence in contrast with a Posnerian interpretation of purely strategic privacy decision-making, and challenges the conclusion that experimental participants' willingness to share data for small rewards betrays lack of demand for privacy tout court. First, an extensive body of work has uncovered numerous hurdles—not just asymmetric information, but also bounded rationality and an array of cognitive heuristics and behavioral biases—that influence (and to some degree impair) strategic privacy decision-making in the marketplace (Acquisti, Brandimarte, and Loewenstein 2015; Acquisti et al. 2020). We will go back to that literature in Section 2.2.3. Second, behavioral research has actually provided clear evidence of extensive privacy-seeking behavior, both online and offline. Writing two decades before the Internet, Altman (1975, 1976) noted that privacy-regulating behaviors are common and sometimes instinctual. Boundary regulation implies a "continual adjustment and readjustment as new situations emerge" (1976, p. 23), with people implementing "desired levels of privacy by behavioral mechanisms such as verbal and paraverbal behavior, nonverbal use of the body, environmental behaviors and cultural norms and customs" (p. 17). Those behaviors may be invisible to us economists merely because they escape our definitions of privacy. Ordinary examples from our daily lives abound offline (we lower our voice or change topic when a third party approaches as we are engaged in an intimate conversation with someone; we step aside from a group of friends when we get the call from the doctor's office with the results of a test) but also online (we alternate between different email accounts or online personae to separate personal from professional spheres; we pick privacy settings to manage the visibility of our social media posts). (See Acquisti et al 2020, from which these examples are taken.) Actual studies (including self-report surveys, observational field works, and experiments) complement the anecdotal observations. For instance: a majority (58%) of social network site users surveyed by Madden (2012) had restricted access to their profiles; only 22% of CMU Facebook users publicly shared their date of birth in 2009 (down from 86% in 2005; Stutzman, Gross, and Acquisti 2013); 50% of participants in an experiment were unwilling to exchange a \$10 anonymous gift card for a \$12 trackable one (Acquisti, John, and Loewenstein 2013); following Apple's transition to the App Tracking Transparency framework

(ATT) in 2021, which imposed an opt-in tracking framework for apps on the Apple ecosystem, an overwhelming share of iOS users opted *not* to be tracked;⁹ and a substantial proportion of Internet users worldwide use adblockers as tools to block unwanted ads from popping up on their browsers (the proportion varies from study to study, from 27% to close to 50%).¹⁰ In fact, a recent study of the “reverse” privacy paradox (the investigation of privacy-seeking behavior among individuals who claim privacy to be of little importance to them), found that engagement in a broad array of privacy behaviors was very common in a US-based online sample of 255 participants. The vast majority of participants reported having engaged in most of the privacy behaviors randomly picked from a list and presented to them. In fact, even a majority of those participants who had claimed privacy not to be particularly important to them had engaged in those privacy-protective behaviors (Colnago, Cranor, and Acquisti 2023).

The empirical behavioral evidence may thus suggest that, contra the notion of digital denizens doing little to protect their privacy, consumers engage in privacy management all the time—that is, they continuously, and often without noticing, make decisions to regulate their degree of openness with others. This does not mean that they want to *protect* their data *every time* (Acquisti et al. 2020). Of course they do not: privacy, from an Altmanian perspective, is about dynamically seeking both openness and closeness, depending on context. In fact, and contrary to the notion of privacy as a modern invention, substantial multidisciplinary research (from history, anthropology, and ethnography, as well as ethology) provides evidence that privacy-regulating behaviors may be a universal trait of human societies across space and time. Such historical universality may be explained by an intriguing conjecture: there may be evolutionary roots to modern privacy concerns (Acquisti, Brandimarte, and Hancock 2022). The ability to detect through our senses the presence of others in our physical space and to recognize friend from stranger or foe and react accordingly provides a clear evolutionary advantage. Over time, as human cognition evolved, so did human ability to negotiate the boundaries between self and others for self-interest: to avoid threats and leverage opportunities. Thus an evolutionary account of privacy can explain the remarkable diversity of dimensions (and definitions) of privacy across time and cultures (as Altman 1977 noted, privacy is simultaneously culturally universal and culturally specific) and can highlight the deep link, now as in our distant past, between the need for security and the drive toward privacy - or, to go back to economic terminology, our *demand for privacy*.

A fourth difference I want to highlight derives from the prior three, and pertains to contrasting stances over privacy regulation. By and large, in other social sciences and in computer science, the value of privacy is often normatively (for economists, perhaps, paternalistically) assumed; strengths and weaknesses of different forms of protection are discussed; and among them, regulation is commonly accepted as a legitimate tool for policy intervention. In contrast, mainstream economic analysis has often been skeptical of or outright averse to privacy regulation (again, exceptions exist: see, for instance, Becker 1980, or Arrieta-Ibarra et al. 2018). At the very outset of the field of research, Posner (1981) lamented “the rash of recent privacy legislation and the high level of public as well as scholarly concern with privacy” (p. 408). A little

⁹ See <https://www.macrumors.com/2021/05/07/most-iphone-users-app-tracking-opt-out/>.

¹⁰ See <https://www.insiderintelligence.com/content/ad-blocking-growth-is-slowing-down-but-not-going-away>.

less than two decades later, Varian (1996) warned that as privacy was becoming a very contentious public policy issue, Congress may “rush into legislation without due consideration of the options. In particular, a poorly thought-out legislative solution would likely result in a very rigid framework that assigns individuals additional rights with respect to information about themselves, but does not allow for ways to sell such property rights in exchange for other considerations” (p. 108 of the 2009 edition). Roughly another 20 years later, in an exceptionally balanced piece, Jin and Stivers (2017) considered several tools and interventions available to policymakers interested in privacy, including educating consumers, voluntary or mandatory disclosures, and minimum quality standards determining how firms should collect, store, use and share consumer data. Although they did not endorse or dismiss any of them, they contrasted interventions that focus on privacy processes, which ensure that “consumers and sellers have the tools to exercise appropriate control on the process” and “should help bolster a *healthy market to facilitate and honor their choice of privacy*” (emphasis added, as later in this manuscript I will get to issue of whether policy interventions such as informational or educational campaigns can in fact assist consumer privacy choice), to “a more paternalistic approach that attempts to determine consumer preferences on privacy outcomes and directly impose that determination on the market.” They also observed that a policymaking body would have such a variety of tools to apply “[o]nce it has decided that a market failure exists and it is likely to cause *net harm* to consumers” (p. 21)—that is, only once economic damage has been established (emphasis added, again as I will go back to the concept of net harm—and whether it can be calculated—in Section 2.2.2). These analytical concerns are reflected in the empirical literature. Echoing Posner’s skepticism toward regulatory interventions, a large share of empirical economic research on privacy has focused on documenting the costs and inefficiencies caused by protection of personal information and privacy regulation (see Section 2.2.1).

Different training and ideological differences can explain in part the gap between economists’ and other scholars’ stances on the merits of privacy regulation. Yet surely that gap is also driven by differences in how economists and privacy scholars *construe* privacy. The four differences I highlighted in this section are logically interrelated. If privacy is construed mainly in terms of concealment and in terms of individual, locally optimized, decision making, then the abundant evidence of online disclosures will be taken as proof of weak individual preferences for privacy; and if rational behavior in the marketplace accurately captures those preferences, it will follow that privacy regulation is unnecessary at best and deleterious at worst. If, instead, privacy is more than concealment and pertains to more than information, then evidence of public disclosures will not be taken as proof that individuals do not care for privacy; in fact, they do, but behavioral hurdles and economic barriers make it hard for them to achieve the privacy they desire in the digital marketplace; hence regulation will be needed to allow individuals to manage their privacy in a world of endemic information asymmetry and systemic power imbalances. I expand on this in Section 2.2.3.

2.2 Unintended Consequences

The success of the economics of privacy as a research field was built in part on a narrow but analytically rigorous focus which pioneers such as Posner and Stigler proposed. That approach

deviates from much of other social sciences' theorizing on privacy. In this section I discuss the unintended consequences of that deviation. Because as economists we sidestep the richness of the multiform dimensions of privacy in the literature outside economics, we end up spending more time focusing on the trees (informational costs) than the forest (the profound ramifications of the evolution of privacy boundaries in our digital societies). In doing so, we insulate ourselves from an array of empirical research questions that go beyond the study of the impact of data protection (Section 2.2.1), from the evidence of widespread consumer privacy harm (Section 2.2.2), and from the implications of privacy behavioral research (Section 2.2.3).

2.2.1 The Disconnect Between Empirical and Theoretical Economic Privacy Research

A first consequence of the narrow economic view of privacy is the disproportionate attention that empirical works have paid to one particular research question. While several exceptions exist (I offer examples below), a common focus of empirical research in this field has been the quantification of economic inefficiencies and costs arising from privacy regulation: from reducing the impact of online ads on hypothetical purchase intentions (Goldfarb and Tucker 2011) to decreasing the speed of adoption of electronic medical records and technologies that can save infants' lives (Miller and Tucker 2011) to reducing ecommerce spending (Goldberg, Johnson, and Shriver 2023)—just to name a few. Individually, these and many other studies are rigorous. In the aggregate, they suggest a disconnect between the dominant empirical analysis and the theoretical privacy economics literature.

The theoretical privacy literature has repeatedly highlighted highly nuanced economic effects of both information protection and information sharing. It has demonstrated over and over again (see a review in Acquisti et al. 2016) that both at the individual level (that is, in terms of individual welfare) and at the societal level (aggregate welfare), privacy protection can be either welfare-decreasing or welfare-enhancing, depending on context. The nuanced effects in terms of individual welfare are the easiest to illustrate intuitively using simple economic theory: Varian (1996) had already pointed out that *not* sharing personal data could both benefit the consumer (when that data was her reservation price) and harm her (when that data was her product preferences). Further, as Noam (1996) observed (and as we noted in Section 1), privacy is a domain where the interests and rights of different parties collide. Thus there is no reason to expect *ex ante* that the interests of both data subjects and data holders will align, nor that the degree of privacy in the market will be optimal for both parties. There is no way (aside, perhaps and sometimes, from privacy-enhancing technologies; see Section 3) to avoid certain trade-offs between data subjects and data holders. If privacy is redistributive, as Posner (1981) proposed, so is the *lack* of privacy (Acquisti et al. 2020).

The theoretical argument that illustrates how the effects of privacy on aggregate welfare may be similarly nuanced is less intuitive, because it can take multiple forms. Several theoretical analyses even show, for instance, that *lack* of privacy protection can *decrease* aggregate welfare. They range from Hirshleifer's (1971) classic argument about private (not necessarily personal) information (the private benefits of information acquisition may outweigh its social benefit; in a pure exchange setting, information may have no social value as it merely results in

a redistribution of wealth; thus, economic agents may overinvest in private information acquisition), to Hermalin and Katz's (2006) ex ante vs. ex post trade efficiency argument (under which the provision of privacy can create welfare-increasing equilibria that otherwise would be destroyed). One illustration of Hermalin and Katz's argument appears prescient today: "[f]or example, absent the ability to keep information confidential, people may not collect information about themselves (e.g., individuals might forgo AIDS testing if disclosure were mandatory), resulting in unintended adverse consequences" (p. 212). Compare this example to the results in Derksen et al. (2022): HIV patients may dodge tracing precisely because of their (often justified) fear that medical conditions will not be kept confidential. Credible assurances of privacy protection may induce patients to consent to tracing, thereby improving both individual and societal well-being.

With—again—important exceptions (such as Marthews and Tucker 2017; Neumann, Tucker, and Whitfield 2019; Buckman, Adjerid, and Tucker 2022), these theoretical nuances rarely surface in empirical works. Even when they do, the economist's skeptical stance towards regulation percolates all the way up to how we frame our results for the public; for instance, the careful study by Buckman et al. (2022) I just cited—which found that privacy protection can *increase* demand for COVID-19 vaccines—was titled "Privacy Regulation and *Barriers* to Public Health" (emphasis added).

One possible explanation for the divide between empirical vs theoretical privacy economic literatures is simple: the empirical literature has tested all sorts of theoretical predictions, but found support only for those which highlight the costs of data protection; the costs *are* there, and those results get published. An alternative explanation is self-selection in how we pick research questions and dependent variables (metrics) to investigate. For various reasons, we tend to pick questions that focus on the costs of regulation—and, often, we find evidence for those costs, since we rely on short-term metrics most likely to capture them. Those reasons may include training, mindset, exogenous events (the enactment of privacy regulations creating favorable conditions for field experiments), as well as researchers' cost-benefit analysis, based on data availability, accessibility, and publishability: it is hard to conduct rigorous empirical investigations of the impact of privacy regulation even on relatively available short-term market metrics (for instance: venture capital investments following the enactment of the *General Data Protection Regulation* (GDPR): Jia, Jin, and Wagman 2021; or app developers' monetization strategies following the introduction of Apple ATT: Kesler 2022; and so forth); it is even *harder* to look at the long-term ramifications of those regulations on more diverse metrics, including possibly beneficial effects—not because the latter ramifications do not exist, but because they are much more difficult to quantify (they may be less tangible; the needed variables may not be readily available from corporate databases) and to causally link to the regulation itself (as those ramifications may manifest progressively in the longer term; I delve deeper into these challenges in Section 2.2.2.) Ultimately, our scholarly drive towards robust identification (which these papers often address with cleverness and rigor) shrinks the space of admissible research questions that can be addressed with sufficient precision to withstand the exacting peer-reviewing process. And given that economic journals are not usually averse to results exposing the unintended effects of regulation, the researcher's cost-benefit calculus can ultimately steer

our choice of research questions. The result is a body of works individually rigorous but collectively incomplete.

2.2.2 The Economic Paradox of Privacy Harm and the Aggregation Problem

A corollary of empirical economic research's focus on the costs of privacy protection—and a second consequence of the narrow economic theorizing of privacy—is the sidestepping of evidence of an extensive amount of consumer and societal privacy harm.

As noted in Section 2.1, the economics of privacy has predominantly focused on informational issues. Accordingly, the literature has concentrated on a limited subset of harms associated with personal *data* and its regulation. For instance, the modeling literature has tended to associate consumer privacy harm with price or product discrimination arising from the tracking of consumer preferences (as in Taylor 2004 or Acquisti and Varian 2005), or with an abstract individual “taste for” privacy, which typically captures an individual's preferences concerning the amount of her personal information available to others (Farrell 2012). As noted, the empirical literature too—with notable exceptions—has tended to focus on measuring data-related harms such as identity theft or the economic impact of regulatory protection of personal information. Because of this, many typologies of consumer privacy harm have been sidestepped by economic research. In fact, the very existence of consumer concerns over privacy has been sometimes a source of explicit *bewilderment* in our field. Consider Posner (1978): “the privacy legislation movement remains a puzzle from the economic standpoint.” Consider, again, Posner (1981): “why people should want to suppress such facts is mysterious from an economic standpoint” (referring to publicizing facts that have no possible value to potential transacting partners).¹¹ And consider, more recently, Wickelgren (2015): “[w]hile concerns about privacy and the collection of consumer information are becoming ubiquitous, they are raised in a fashion that is puzzling to an economist. That is, they typically do not explain what potential market failures may exist that would lead the market not to provide the optimal amount of privacy when consumers use internet services such as search engines or shopping platforms.” To be fair, theoretical work (e.g., Becker 1980, Hermalin and Katz 2006; Farrell 2012) did acknowledge the existence of distinct consumer preferences for privacy as an “intermediate” good (whose value is instrumental—e.g., protecting privacy to avoid identity theft) and as a “final” good (whose value is intrinsic—e.g., protecting privacy because of personal taste). But empirical estimates of the vast array and diversity of harm discussed at length in the legal privacy scholarship (Calo 2011; Citron and Solove 2022) are lacking within economics. Empirical evidence in adjacent fields such as communication research or human-computer interaction *has* repeatedly highlighted consumer privacy concerns with several commercial data practices (McDonald and

¹¹ A law scholar alerted me of a 2004 decision signed by Posner, as Circuit Judge, in *Northwestern Memorial Hospital V. Ashcroft*, written over 20 years after his seminal economic analysis of privacy. The decision highlights the value of (medical) privacy: “Even if there were no possibility that a patient's identity might be learned from a redacted medical record, there would be an invasion of privacy.” The same scholar, however, added: “Some Posner scholarship after his Northwestern Hospital decision returns to his privacy-skepticism” (personal communication with the author).

Cranor 2010) and systemic gaps between their privacy expectations and those practices (Rao et al. 2016, Turow, Hennessy and Draper 2018). But those concerns and those gaps are rarely recognized as economic harm. In the United States, Courts have increasingly expressed skepticism towards the notion that individuals who merely felt that their privacy was violated—but only suffered injuries which were difficult to quantify—should be able to sue (see Strahilevitz and Liu 2022).

Once we look at privacy research outside economics, we realize that the paradox of privacy harms is that their measurement is hard not because of their rarity but for the opposite reason: privacy harms are ubiquitous but diverse in form, heterogeneous in likelihood, and varying in magnitude and length. These disparate and context-dependent embodiments of harm make it hard to quantify or even just conceptualize privacy damages into a single intuitive metric. We have referred to this as the *aggregation problem* (Acquisti et al. 2020). Harms associated with misuses of personal data include both those immediately recognizable as economic costs and those with less directly quantifiable (yet no less important) repercussions, such as physical harm, reputational harm, psychological harm, autonomy harm, discrimination harm, and relationship harm (Citron and Solove 2022). Under each of these categories, numerous distinct sub-instances of harm can be defined: from identity theft to price discrimination, from attention and time waste to chilling effects, from hiring discrimination to filter bubbles narrowing individual choice, from stigma and psychological harm to rare but catastrophic physical consequences, and more. Commercial surveillance (Zuboff 2015) practices that increase the amount of consumer data collected and shared with third parties—often without individuals' knowledge and consent—ultimately increase the stochastic risk that any one of those myriad possible harms may occur. Therefore, while the likelihood of any individual type of harm occurring may be low, the typologies of possible harms are so many that surveillance practices ultimately elevate the statistical expected cost of commercial surveillance for each consumer and for the aggregate of consumers. And yet that expected cost remains hard to quantify (for scholars, policymakers, and the consumer herself) because of the aggregation problem.

Consider the following examples of the myriad scenarios in which the collection of consumer data has tangible, significant, and far-reaching ramifications which remain challenging to capture in economic analysis.

Scenario 1: every time a person visits a website, the time it takes for its content to load is extended by the plethora of trackers that collect information about the visitor and pass it to other third parties for the purpose of online advertising. This happens on the vast majority of websites. This transaction cost is small at the individual visit level.¹² Aggregated across multiple visits

¹² But not negligible. Borgolte and Feamster (2020) tested how privacy-focused browser extensions for Google Chrome and Mozilla Firefox affect browser performance. In their tests, while using those extensions came at some cost, those costs were offset by performance improvements due to blocking tracking. They write: “[c]ontrary to Google’s claims that extensions which inspect and block requests negatively affect browser performance, we find that a browser with privacy-focused request-modifying extensions performs similar or better on our metrics compared to a browser without extensions” (p. 2275). For instance, they report that extensions that merely block online trackers, such as Disconnect, can reduce actual page-load time by as much as 244ms (median) - nearly a quarter of a second per visited page, per user.

conducted by an individual over time, and across multiple individuals, the aggregated opportunity cost of time lost to trackers is significant. This scenario is an example of a widely common (high likelihood) cost that is minimal at the event-level but remarkable in the aggregate.

Scenario 2: in a handful of cases, American prosecutors “have used text messages and online research as evidence against women facing criminal charges related to the end of their pregnancies.” For instance, in 2017, a Mississippi woman, Lattice Fisher, “was charged with second-degree murder after a failed pregnancy [...] Prosecutors drew heavily on Fisher’s search history. Notably, local reporting claims the police found record of these searches from Fisher’s own phone rather than through Google itself.”¹³ Following the US Supreme Court’s overturn of *Roe v. Wade* with its 2022 *Dobbs v Jackson Women’s Health Organization* decision, concerns have grown over the way police agencies may use search, browsing, or app data against women who merely tried to learn about abortion (Ms. Fisher’s case was later dismissed, but only after she had spent time in jail).¹⁴ This is an example of an event with low probability of occurrence but major individual consequences.

Scenario 3: in September 2018, a UN report highlighted the role of social media in fomenting hatred and ultimately genocidal violence (including mass killings, rapes, and destruction) in Myanmar.¹⁵ The report called out Facebook as “a useful instrument for those seeking to spread hate” (p. 14): the Myanmar military had used Facebook systematically to engage in propaganda against the Rohingya people. Facebook itself, through an independent report it commissioned to the BSR (Business for Social Responsibility) admitted its role in not “doing enough to help prevent our platform from being used to foment division and incite offline violence.”¹⁶ It is important to point out the central role the tracking of personal data by social media platforms plays in these and similar societal dynamics. That role is central not merely because social media relies on the monetization of personal information for its sustainment (for instance, via targeted advertising); but also because personal information is critical to foster engagement. Algorithms use personal data to select which information to show to which users to increase the amount of time they spend using the services, and get exposed to ads. And those algorithms may be blind to whether they are encouraging a visitor to watch one more video about their favorite football team—or they are riling her up with rage against the purported misdeeds of another group of people. This is an example of a very common occurrence (algorithmic targeting) contributing to (among other things) exceedingly rare events with catastrophic individual and societal consequences (genocidal violence).

Scenario 4: Bradshaw and Howard (2018) found evidence of organized social media manipulation campaigns in 48 countries, with at least one party or government agency in each of the analyzed countries using social media to manipulate domestic public opinion domestically including through disinformation campaigns. As in the Myanmar case, personal data play a

¹³ See <https://www.theverge.com/23185081/abortion-data-privacy-roe-v-wade-dobbs-surveillance-period-tracking>.

¹⁴ See <https://www.pregnancyjusticeus.org/victory-for-lattice-fisher-in-mississippi/>.

¹⁵ See https://www.ohchr.org/Documents/HRBodies/HRCouncil/FFM-Myanmar/A_HRC_39_64.pdf.

¹⁶ A. Warofka, “An Independent Assessment of the Human Rights Impact of Facebook in Myanmar,” Facebook (2018; revised 2020); <https://about.fb.com/news/2018/11/myanmar-hria/>.

central role in these operations, especially via misinformation designed to target and appeal to specific groups. And yet, while social media *may* sway small but ultimately key portions of voters in very close elections (Aral and Eckles 2019), it may be impossible to conclude definitively whether and when an election was won or lost due to how unknowing voters' data was used to target them.¹⁷ In fact, the very ability of so-called filter bubbles to significantly affect downstream societal dynamics has been subject of debate (Bruns 2021). Considering the far-reaching ramifications (both social and economic) of a nation voting in one leader over the others (or social media platforms amplifying already-occurring dynamics of polarization), researchers and policymakers thus face a paradox and a challenge. The paradox is that data-driven online campaigns may have downstream effects on the citizenry that are potentially staggering, yet for which it is impossible to rigorously demonstrate and precisely estimate a causal relationship. The challenge is that the more we attempt to decrease the probability of a Type I error in investigating those relationships and in guiding policy, the more we risk making a Type II error: dismissing the potentially far-reaching social ramifications of the loss of privacy.*

To emphasize complexity and heterogeneity, the four selected scenarios vary in likelihood, magnitude, and typology of privacy harm. And yet they are mere examples from a broader and potentially unbounded set. Countless other scenarios and alternative downstream harms may exist, because, once collected, the boundaries of usage of personal information are undefinable and unpredictable.

Lin (2022) has estimated and compared instrumental and intrinsic preferences (valuations) for privacy—a distinction similar to Farrell's analysis of privacy as an intermediate and as a final good. It is important to note that the distinction between instrumental and intrinsic preferences for privacy is different from the measurement of different typologies of realized consumer privacy harm we are considering here. Such harm is stochastically realized and unpredictable *ex ante*. Thus it is independent of both a consumer's intrinsic preference for privacy and—due to information asymmetries—of her expected economic trade-offs from sharing or protecting data as well. For example, a consumer may bear high material costs from identity theft regardless of how privacy sensitive she is, and independently of whether she expects her identity to be stolen.

While the *ex post* realization of consumption utility from any economic good may deviate greatly from its *ex ante* anticipation and consumer expectation (the costly car the consumer purchases could turn out to be a lemon), the case of privacy is unique. Data, unlike physical goods, can be nonrival (Jones and Tonetti 2020) and nonexclusive, and once revealed is subject to repeated, potentially unending secondary use. Varian (1996) first observed that widespread secondary use of digital data could give rise to externalities. Individuals rarely know or predict the many

¹⁷ At least in the case of the 2016 EU referendum campaign in the UK and the Cambridge Analytica scandal, the letter by the UK Information Commissioner on the investigation into use of personal information and political influence ultimately found that Cambridge Analytica was "not involved in the EU referendum campaign in the UK" (p. 2). The Commission however also confirmed the existence of "systemic vulnerabilities in our democratic systems" associated with new tracking and targeting technologies (see https://ico.org.uk/media/action-weve-taken/2618383/20201002_ico-ed-l-rtl-0181_to-julian-knight-mp.pdf).

* This challenge is underscored by recent and seemingly contrasting results of studies investigating the impact of social media (and Facebook specifically) on variables such as political polarization or subjective well-being. Contrast Nyhan et al (2023) and Guess et al (2023), who find an amplifying but not polarizing effect of exposure to like-minded sources or reshares on Facebook, to Allcott, Braghieri, Eichmeyer, and Gentzkow (2020), who find that deactivating Facebook for the four weeks before the 2018 US midterm election did reduce political polarization. Or contrast Vuorre and Przybylski (2023), who do not find an association between Facebook use and measures of subjective well-being (using observational data) to, again, Allcott et al (2020), who find a negative association (using a field experiment).

possible secondary uses of their data or their consequences. Examples in the literature abound: now and again, new ways to collect and use personal information are discovered, and users' expectations regarding the privacy of their data are often distant from reality (for instance, see Liu, Gummadi, Krishnamurthy, and Mislove 2011). Information asymmetries are systemic and endemic in the privacy domain. And since the value of data—and thus of privacy—can often be determined only *ex post* (that is, based on the context in which information is used),¹⁸ even a consumer who knowingly engages in a data transaction with another party will ultimately face trade-offs she is not able to predict, account for, or control as a rational economic agent.

Those data externalities may be both negative and positive. But the peculiar (among other economic goods) combination of lack of consumer awareness regarding data uses and lack of control over those uses are precisely what makes it impossible for consumers to make optimizing decisions, reducing the risk of negative externalities while increasing the probability of positive ones (choosing consumption levels of privacy, so to say, to match its marginal costs to its marginal benefits). One cannot optimize for something they neither know nor control.

The context-dependent nature of privacy harm and its *ex-post*-determined trade-offs also raise serious questions over the ability of data markets to fairly capture the value of privacy. At worst, they may make it hopeless to attempt to aggregate privacy net harm into a single economic estimate. Like the consumer, the regulator thus faces the challenge of comparing the social marginal costs to the social marginal benefits of personal data. But the empirical privacy literature stops short of helping the regulator. None of the hurdles we discussed—the aggregation problem, the unbounded set of data usages and consequences, and the entanglement of positive and negative data externalities—can reasonably support the conclusion that consumer privacy losses have no harmful effects on consumer welfare aside from subjective concerns. The economics of privacy has to a great extent sidestepped the evidence of consumer privacy harm. Because of that, we measure the tip of the iceberg and remain unfamiliar with its mass underwater.

2.2.3 Implications Arising from the Behavioral Literature

A third consequence of the narrow economic theorizing of privacy is a misapprehension of the implications of several decades of behavioral privacy research.

As we noted earlier in this section, mainstream economics and behaviorally-focused research have interpreted differently the results of empirical studies of consumer demand for privacy. Mainstream research, following Posner's mold, tends to believe in a process of rational decision-making. Under this account, consumers' online behaviors adequately capture their

¹⁸ Leakages of jogging patterns from your exercise app may alternately lead to your learning new tips and techniques, receiving undesired advertising, or—if you are a military officer in a war zone—getting killed (<https://meduza.io/en/news/2023/07/11/killed-former-submarine-commander-in-krasnodar-could-have-been-tracked-by-running-app>).

demand for privacy. The hurdles consumers face in making privacy choices (especially asymmetric information) are at times acknowledged by careful scholars (see, for instance, Jin and Stevers 2017), but informational and educational interventions are presented as viable strategies to assist privacy-conscious consumers.

Conversely, behaviorally-focused research tends to highlight how those hurdles distort revealed preferences for privacy in the marketplace. According to this account, informational, behavioral, and economic hurdles, far from being sidenotes or exceptions, are ubiquitous, systemic, and central in consumer choice. Hence they make consumers' desired degrees of privacy unattainable through market interactions. Ultimately, no amount of informational or educational intervention may remedy those systemic barriers.

To understand why, let us consider those hurdles. Purely information hurdles (such as asymmetric information) have just been considered near the end of Section 2.2.2. From a behavioral perspective, educational and informational interventions do not necessarily ameliorate those informational hurdles and thus consumer privacy decision making. First, the behavioral literature also suggests that education and transparency, by themselves, are ineffective—they may be necessary but not sufficient tools for privacy management. Notice and consent regimes do not even resolve the basic problem of information asymmetry: they are exorbitantly costly for end-users (McDonald and Cranor 2008), unhelpfully ambiguous and therefore unactionable (Reidenberg et al. 2015),¹⁹ and crash under the weight of both the myriad privacy notifications, options, and requests consumers are inundated with daily²⁰ and our innate bounded rationality. Second, educational and informational interventions crash against a second set of hurdles: a vast array of cognitive and behavioral factors that can affect and impair privacy decision-making (Acquisti et al. 2015, 2020), and which in fact can be exploited by platforms and services providers via so-called dark patterns (Acquisti et al. 2017): whoever controls the user interface controls the architecture of choice.

Drawing attention to those behavioral factors is far from suggesting that consumer privacy behavior is irrational, or that privacy choices are erratic and unaffected by preferences, incentives, and calculus. Rather, it means emphasizing that privacy decision-making deviates in systematic ways from the theoretical prediction of rational choice models, which assume complete information, stable preferences, and procedural invariance—all assumptions the empirical privacy literature has shown untenable (Rao et al. 2016; Acquisti et al. 2013; Tomaino, Wertenbroch, and Walters 2021). As we noted elsewhere (Acquisti et al. 2015, 2020), privacy decision-making (as decision-making in general) is rather the result of both deliberative (utility-maximizing) and behavioral factors.

¹⁹ In fact, even when consumers do read privacy notices, their interpretation of what actual data policies those notices entail seems “depend more on their preexisting expectations” than on the terms of the notices themselves (Strahilevitz and Kugler 2016, p. S71).

²⁰ Skiera, Miller, and Jin (2022) find that if a user were to make all possible decisions regarding the provision of permission for data processing under the GDPR for each new publisher she visits in a day, she would spend 79.13 minutes per day in “decision time.” See also Cooper (2023).

In fact, the evolutionary account of privacy concerns we have presented in Section 2.1 offers a unifying explanation for the various informational and behavioral hurdles we have chronicled here. In the offline world, privacy management is often instinctual, almost natural (which does not imply, however, that one can always achieve the privacy they desire). Online, privacy management is more arduous because of an evolutionary *mismatch* (Pani 2000): we lack the cues humans have evolved to rely on to manage the boundaries of public and private, to detect the presence of others and react accordingly. As we travel on a crowded train, we quickly sense another person's peeking at the documents open on our screen; as we walk in a street, we notice the steps of someone following us too closely. On the Internet, we do not *see* or *hear* Facebook or Google tracking us across all sorts of digital domains. Notice and consent mechanisms—as well as educational or informational interventions—fail because they do not account for the underlying nature of consumer privacy decision-making. Worse, they amount to exercises in consumer *responsibilization*—that is, asking consumers to take charge of a problem they did not create and cannot really control. And they do little to solve the worsening problem of user interfaces designed to nudge consumers toward more engagement and self-disclosure.

Lastly, other hurdles arise from the “supply side” of privacy (we have considered them extensively in Acquisti et al. 2016). Economic barriers make it overly costly for consumers to comprehensively manage their digital privacy, and they often render privacy options entirely inaccessible. These include lack of viable market alternatives (or alternatives being exceedingly onerous); switching costs, adoption costs, privacy externalities, and so forth.

Informational, behavioral, and economic hurdles combine to cripple consumers' ability to manage online privacy. In Altman's terms, they render *achieved* privacy outcomes different from *desired* ones, thus justifying calls for policymakers' intervention.

Related to this discussion is a specific and contentious stream of behavioral work that has been the object of particular misapprehensions and thus confusion about the implications of empirical research: the privacy “paradox.” The paradox is the purported gap or dichotomy between privacy mental states (such as preferences, attitudes, or even intentions, often reflecting a claimed desire for privacy) and actual behaviors (seemingly reflecting a carelessness towards privacy). Few other areas of privacy research have attracted as much attention and caused as much disagreement as the privacy paradox: Is it real, or is it a myth (Solove 2021)? In recent works outside the economic domain (Colnago et al. 2023; Acquisti et al. 2020), we have argued that much of the disagreement over the paradox of privacy has been caused by conceptual confusions. I summarize here a few key points that may be of relevance to the economic debate. A first source of confusion is that the very term “paradox” is interpreted differently by different scholars in the field.²¹ This leads to disagreements over the *paradoxical* (or not) nature

²¹ As noted in Acquisti et al. (2020), the term “paradox” has two similar but subtly contrasting meanings: a “self-contradictory statement that at first seems true” (Merriam-Webster), but also a “seemingly contradictory” statement that is “perhaps true.” The dichotomy between stated mental states (such as preferences or intentions) and behaviors is the (apparent) contradiction. Some scholars appear to look at the dichotomy through the lens of the first definition: they search for explanations of that dichotomy, and when they find them, they conclude that there is no self-contradiction, and thus also no paradox (see, for instance, Solove 2021). Other scholars appear to look at the

of a possible mental states/behaviors gap that are entirely lexicological—if a paradox has an explanation, is it still paradoxical? Opinions vary—and thus have little bearing on the actual empirical comparison of those mental states and behaviors. A second and more consequential source of confusion is the seemingly implicit assumption in much of the work in this field that the question “do privacy mental states match behaviors?” can be answered broadly and conclusively in static, binary terms: yes or no. This, of course, is folly: answering that question in such terms would require believing that attitudes must either always match behaviors or never do. Whereas everything about privacy (including decision-making) is dynamic and contextual. Thus it is more plausible to expect that privacy attitudes, preferences, and mental states will sometimes predict and match behaviors (Dienlin and Trepte 2015) and sometimes will not (Norberg, Horne, and Horne 2007). The gap between privacy mental states and behaviors is therefore neither a myth, nor is it always guaranteed.

This brings us back to the issue of what policy implications to draw from the evidence that sometimes (but not always) a gap will exist between mental states and behaviors, and what implications to draw from the behavioral privacy literature at large. The privacy paradox has acted as a Rorschach test, to which people assign the most diverse interpretations based on their own assumptions and from which they thus draw the most diverse policy conclusions. One conclusion (with which I disagree) is that the privacy paradox literature demonstrates that people do not really care about privacy or do not really know what they want, and therefore no public intervention is needed, other than perhaps some informational intervention. A different conclusion (with which I agree) is that the existence of a gap between mental states and market choices reflects precisely those economic and behavioral hurdles that we have identified in this section, which justify or may even require public policy intervention.

2.3 The Inversion of the Overton Window of Privacy Debate

I have highlighted both successes and unintended consequences of the narrow theorizing of privacy embraced by much contemporary economic research. In concluding this section, I consider the ultimate (if potential) repercussion that embrace may produce: economic arguments progressively crowding out non-economic arguments in the public policy debate around privacy. If this risk were to materialize (and, I argue, there are signs of that happening), it would represent a remarkable inversion of the “Overton window” of legitimate policy discourse around privacy.

In the 1990s, Joseph Overton—a political scientist at the Mackinac Center for Public Policy—argued that politicians are constrained in their support of policies by a “window” of acceptability,

dichotomy through the lens of the second definition, which puts the emphasis on the fact that statements that are seemingly in contradiction could in fact be simultaneously correct. For the latter scholars, it's the dichotomy that is paradoxical, even though it can be explained; for them, the fact that dichotomies between privacy attitudes and behaviors can be explained does not imply that the underlying dichotomies do not in fact exist. Ultimately, focusing on the “paradoxical” nature of the gap (that is, focusing on whether the gap is paradoxical, or is a myth) no longer seems productive, because the disagreement over this point is more driven by grammar than empirical evidence. It would be more fruitful to focus on when, whether, and how, behaviors match vs. deviate from mental states.

which includes the policies that, at any given time, a society accepts as legitimate options.²² That window can shrink or expand based on how societal values evolve. A radical or even unthinkable idea can, over time, become popular and thus acceptable and ultimately be embedded in policy. Through a reverse process, a once legitimate and acceptable idea can, over time, become radical and eventually unacceptable.

How does the concept of an Overton window apply to privacy? In Section 2.1, I pointed to scholarly research indicating that the *drive* for privacy is not a modern phenomenon. Evidence suggests it is a universal (albeit ever fluctuating) construct in human cultures across history and geography. The same, however, cannot be said of the notion of privacy as a fundamental human *right*. Construing privacy as a right is a modern development, a process that has panned out progressively and unevenly over time across different cultures (Hixson 1987). Through that process, in the second half of the 20th century the notion of privacy as a fundamental right reached sufficient legitimacy to be ingrained in the principles of an economic organization such as the OECD. In its 1980 *Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, the OECD remarked that privacy protection laws had been introduced in several member countries “to prevent what are considered to be violations of *fundamental human rights*, such as the unlawful storage of personal data, the storage of inaccurate personal data, or the abuse or unauthorised disclosure of such data” (Preface; emphasis added).²³ The *Guidelines* added, “[m]ember countries have a common interest in protecting privacy and individual liberties, and in reconciling fundamental but competing values such as privacy and the free flow of information.”

That process—which saw even economic organizations legitimize privacy as a fundamental human right—may have started reversing in the 21st century. The rise of the economics of privacy has not merely provided a useful analytical complement to values-grounded views of privacy, but may have diminished the currency of notions of privacy as a right by framing data (and privacy) as tradable assets. When Posner (1977, 1978, 1981) outright dismissed attempts to link privacy to broader values such as freedom and autonomy, his contemporaries (Baker 1977; Bloustein 1977; Hirshleifer 1980) recoiled. They balked at the reductionist viewpoint Posner had espoused. Bloustein (1977) wrote: “Posner’s theory is simplistic, not simple, because it accomplishes its objective by avoiding, rather than confronting, complexity. He seduces by reduction, rather than convincing by explanation. The simple analytical elements of the scheme do not add up to the complex whole. His Truth about Privacy turns out to be some truth about one aspect of privacy” (p. 429). Yet Posner’s framework flourished within economics and over time may have influenced public policy. When the OECD in 2013 revisited its 1980 *Guidelines*, the term “fundamental value” had replaced the original “fundamental right.”²⁴ In fact, the term “fundamental human right” was no longer to be found in the revised *Guidelines*. The recognition of a “fundamental right” was no longer explicitly linked to privacy—even though it was explicitly used in reference to *other* rights, such as freedom of speech, freedom of the

²² See <https://www.mackinac.org/7504>.

²³ See <https://www.oecd.org/sti/ieconomy/oecdguidelinesontheProtectionofPrivacyandTransborderFlowsOfPersonalData.htm>

²⁴ See https://www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf

press, and an open and transparent government, which “[p]rivacy rules should also consider” (p. 35). The 2013 revision also replaced the term “danger” (to privacy and individual liberties) with the term “risk” (p. 35), reflecting an increased emphasis on risk assessment. What else had changed? The terms “right” and “economic” appeared thirty-two and seven times, respectively, in the 1980 *Guidelines*. They appeared sixty-one times and forty-eight times, respectively, in the 2013 revision, reflecting both the phenomenal growth of the data economy and the evolution of our priorities in discussing it.

The encroachment of economic considerations in matters of privacy policy was not limited to OECD documents. As the number of lobbyists for the data industry kept growing in Brussels and DC in recent decades, industry-funded think tanks increasingly promoted data-economics arguments against the enactment of privacy regulation. Not coincidentally, references to economic considerations (such as consumers’ right to opt out of sale of their data or businesses’ legitimate interest to process data) and economic factors appeared in regulations such as the *California Consumer Privacy Act* (CCPA) in the US and the GDPR in the European Union. Even the historical 2022 *Rulemaking on Commercial Surveillance* by the Federal Trade Commission included numerous questions aimed at quantifying or estimating the economic dimensions of privacy.

Economist colleagues may disagree with my interpretation of the trends of the privacy debate, and may spot an opposite trend. They may lament—much like Posner four decades ago—regulators’ archaic reliance on values-based normative arguments and their blindness to the soundness and objectivity of economic arguments. Some may even consider what I detect as an emergent unintended consequence to be a very much intended and well needed progression in the policy discussion around privacy. Yet if values-grounded arguments had remained so powerful and persuasive among policymakers, US regulators would eventually have implemented the OECD principles from the 1980s—which stipulate mandatory standards of protection for all personal data—rather than the patchwork of notice and consent approaches still dominant today (and which we have critiqued in prior sections). On the contrary, the influence of economic considerations and industry interests has been evident even in the evolution of drafts of comprehensive European policy interventions such as the *GDPR* (Atikcan and Chalmers 2019; Christou and Rashid 2021). Considering the vast network of organizations lobbying against privacy regulation—as well as the inherent power asymmetry between the concentrated economic interests of large industry players and the diffuse, atomistic interests of uncoordinated individual citizens (Olson 1965; Acquisti et al. 2020)—a once unthinkable scenario now seems possible: the Overton window of acceptable discourse around privacy may be inverting. After a centuries-long evolution in the direction of construing privacy as a fundamental right, the very act of valuing privacy independently of economic evidence may be deemed naive, and eventually radical in some circles. An emerging policy mindset would be that, if there is no easily quantifiable economic harm, then there is no privacy concern worth worrying about. Under such a mindset, policymaking would narrow its focus on what our field has been able to quantify in economic terms—at the risk of discounting harder-to-quantify evidence of privacy harm.

Even nowadays, at economic conferences, I have observed scholars anticipating and preemptively shutting down (in the mold of Posner’s 1981 article) references to freedom or

autonomy, policing the contours of acceptable economic discourse around privacy. Delimiting the contours of the debate is, of course, laudable when our goal is to safeguard rigor in analysis, and when we use the results of our precise but narrow economic observations as complements to the findings of other fields. Delimiting the contours of the debate is instead problematic if we do not exercise similar restraint in also delimiting—carefully and publicly—the *scope* of our contributions—that is, when we use economics as a substitute for other findings to influence public policy and public discourse. Yet such restraint is rarely exercised in our writings. The custom began with Posner. In 1978, he commenced his piece on “An economic theory of privacy” by stating, “I will sidestep the definitional problem by simply noting that *one aspect of privacy is the withholding or concealment of information*” (p. 19; emphasis added). After focusing his analysis on that one aspect, Posner ended the piece on much broader terms: “[i]n the perspective offered by economics and by the common law, the recent legislative emphasis on favoring individual and denigrating corporate and organizational privacy stands revealed as still another example of perverse government regulation of social and economic life” (p. 26). Contemporary economic literature on privacy is not as acerbic, but often follows a similar rhetorical template: the benefits of modern data analytics are espoused at the onset of our articles; the (typically negative) effects of regulation that protects personal information may have on those benefits are then analyzed; performative and typically perfunctory references to privacy’s other dimensions are interjected, sometimes; but then broad, encompassing warnings to regulators (with pleas to consider carefully the unintended consequences of their interventions) are offered as conclusions.

As economists, we are certainly permitted to articulate the implications of our research.²⁵ What we should be wary of is the risk of an intellectual sleight of hand: studying a part (the effects on a subset of directly measurable, hand-picked metrics) but making conclusions for a whole (broad warnings to regulators) that our analyses have barely grazed.

Section 3: Turning the Tables: The Economic Argument for Privacy

A rhetorical template originated with the 1980s economics of privacy literature: limiting the scope of analysis to a particular dimension of privacy, but broadening the implications of that analysis to encompass privacy at large. That template exemplifies a particular way of framing the public debate around privacy. Figure 1 crudely captures key features of that framing. The rest of this section critiques it.

Under the framing that economics has popularized within the public discourse around privacy, a metaphorical scale is weighing two possible outcomes. One outcome is “more” privacy (for instance, regulatory interventions enforce minimum data protection guidelines, privacy-enhancing technologies are deployed, and so forth). The other outcome is “less” privacy and more liberal flows of personal data. The scale measures and compares the benefits to humanity of those two outcomes. Inherent to that framing is the assumption that interventions such as

²⁵ I have done so. When it comes to drawing implications from privacy research, I have found it worth distinguishing two related but distinct questions: should digital privacy be better protected? If so, how? I find the former question harder to resolve in purely economic terms (see Section 3), but have been more sanguine about the latter and thus about articulating the policy implications of available research that addresses it.

regulation aimed at protecting privacy may increase abstract benefits such as freedom or autonomy (which are measured on the right-side pan) but may threaten the more tangible economic benefits from data, such as more free content and services, more innovation, more efficiency, and so forth (which are measured on the left-side pan). Vice versa, refraining from regulating privacy may harm intangibles like freedom and autonomy but may allow more concrete economic benefits to be extracted from data.

The rest of this section argues why this scale—and in fact this way of implicitly or explicitly framing the debate around privacy—is flawed. It is flawed not on abstract moral grounds, but on objective economic grounds. The section argues that uncritically (or unknowingly) internalizing this framing of the debate—as a contest or trade-off between benefits of more data versus the value of more autonomy, dignity, or control—is an erroneous reading of the available scholarly evidence around privacy. Section 3.1 focuses on the left pan. It questions how much we actually know about the allocation of benefits from consumer data, and concludes that we assume a lot but know little. Section 3.2 focuses on the “beam”—the assumption that privacy protection is inherently and inevitably antithetical to the extraction of societal value from data. It rebuffs that assumption and challenges the notion that data protection is inherently welfare-decreasing. Section 3.3 focuses on the right pan. It highlights how little we know about the economic ramifications of privacy invasions. Finally, Section 3.4 proposes alternative ways of framing economic research around privacy, suggesting research questions that are complementary to the current focus on the costs of privacy regulation and proposing a reframing of the economic debate around privacy.

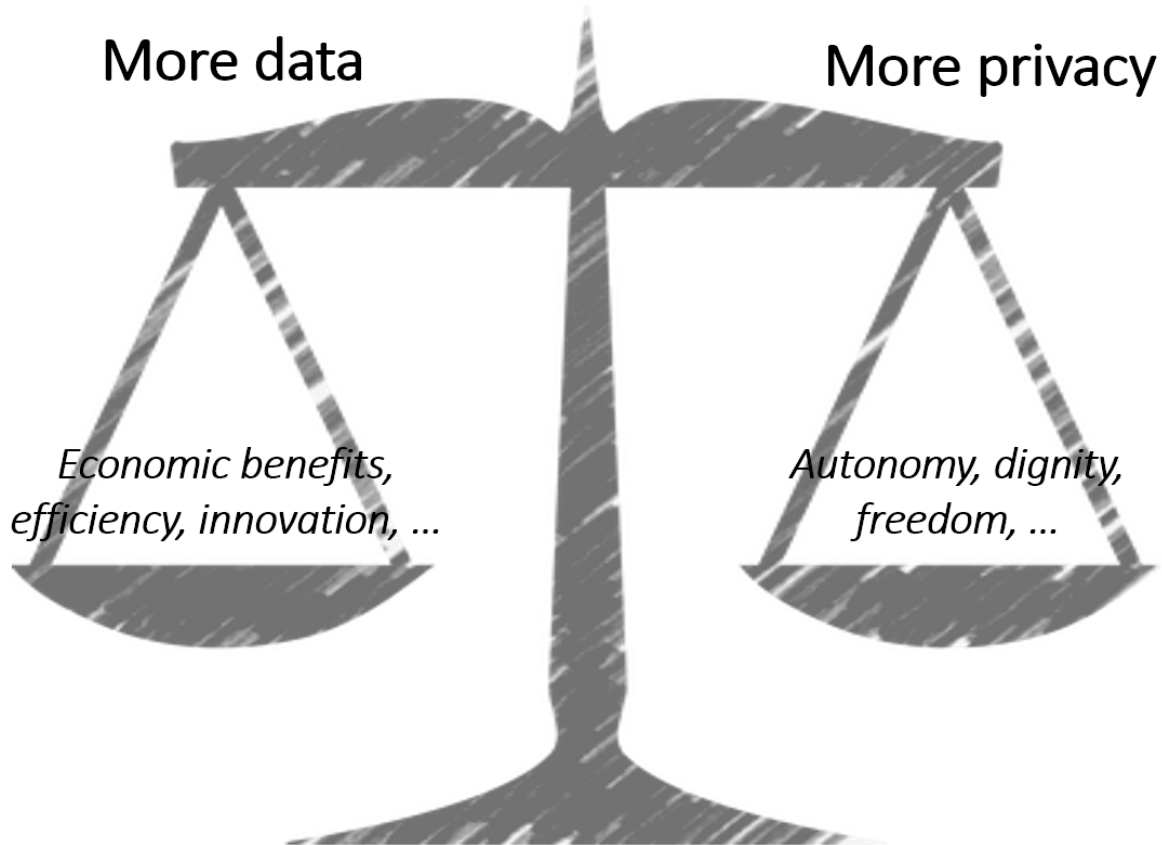


Figure 1

3.1 Missing the Forest for the Trees: What Do We Know about the Allocation of Benefits in the Data Economy?

The left pan of the scale presented in Figure 1 measures the economic benefits that arise from consumer data collection. How much do we actually know about those benefits, and their allocation to different stakeholders, including consumers themselves? That societies can extract value from consumer data is undeniable. But can we separate the snake oil of analytics from its demonstrable gains, and identify the allocation of those benefits?

Extant economic research falls short of these goals. I will focus, as a case study, on the online advertising market. It is not the only sector in which consumer data is tracked and analyzed. However, historically, it has played an outsized role in the process through which the Internet became an architecture of commercial surveillance, and in channeling consumer data into a black box of secondary uses and applications.

A quote from an online advertising executive published in *AdExchanger* (an online magazine related to the online advertising industry) in 2011 captures a widespread way of thinking about the benefits of online advertising, and in particular behaviorally targeted advertising—one of the key innovations in advertising made possible by consumer tracking:

Behavioural targeting is not only good for consumers it's [sic] a rare win for everyone. [...] [It] ensures that ad placements display content that you might be interested in rather than ads that are irrelevant and uninteresting. [...] Advertisers [...] achieve [...] a greater chance of selling the product. Publishers also win as being able to offer behavioral targeting increases the value of the ad placements.²⁶

The notion of behavioral advertising as an economic win-win for multiple stakeholders is consistent with some of the academic literature more critical of regulatory privacy interventions. Figure 2, left side (Frame 1) presents an economic interpretation of that notion. The Figure represents online advertising as a two-sided platform market. Consumers (who visit online publishers, by which I refer to outlets producing various contents and services) want to find merchants to buy from. Merchants (who advertise on the publishers' websites) want to find consumers to sell to. Significant search costs exist on both sides of this market. The data economy intermediaries (companies such as Google, Meta, and other stakeholders in the ecosystem) play the role of matchmakers.²⁷ They use the vast amount of consumer and merchant or product data they collect to facilitate matching between consumers and merchants, via the publishers. By doing so, they reduce search costs on both sides of the market (in particular, for smaller firms trying to reach niche consumer segments), and increase efficiency. Thus they create economic utility (value creation is symbolized by the little arrows coming out of the intermediaries box in the directions of merchants and publishers/consumers). Under this framing, online (behavioral) advertising does create economic win-win for all stakeholders in the market.

²⁶ See <https://www.adexchanger.com/online-advertising/why-is-tracking-good/>.

²⁷ The online advertising ecosystem is, of course, more complex than how Figure 2 depicts it. There are different types of intermediaries, and some intermediaries may also act as publishers and or advertisers. I am abstracting from those details to focus on its key trends.

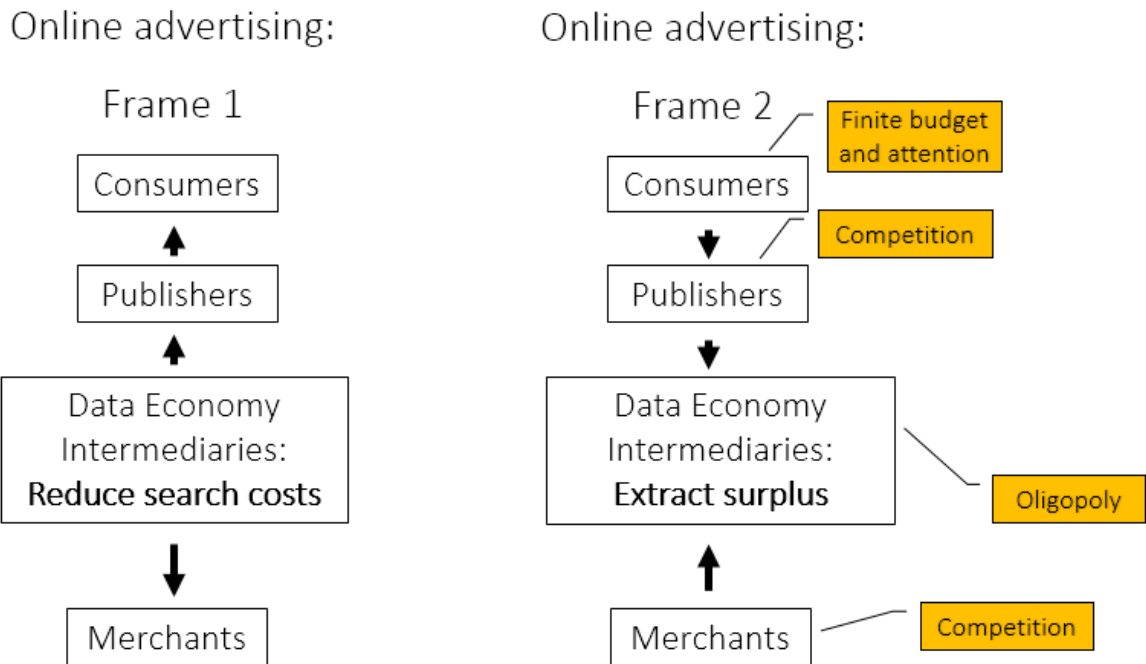


Figure 2

The right side of Figure 2 (Frame 2) presents an alternative economic representation of the same market. The stakeholders are the same. The focus, however, changes from search costs to competition, and from the role of intermediaries in reducing search costs to their ability to extract surplus from both sides of the market. This alternative economic interpretation of the market is equally legitimate, on theoretical grounds, to the economic win-win scenario depicted on the left side, but—as we will see—its conclusions regarding the allocation of benefits from data are different.

Central to Frame 2 is the observation that consumers have finite budgets and finite attention; they cannot pay attention to all the ads shown to them online and cannot purchase all the products advertised to them. Therefore publishers aggressively compete with each other for that limited consumer attention, and merchants compete aggressively for that limited budget. This has several consequences for those stakeholders.

I will consider publishers first. The rise of online advertising (and then of behavioral advertising specifically)²⁸ has acted as a double-edged sword for publishers. First, it has supported the creation of new content dissemination channels and supported new content creators; in doing so, it has increased competition faced by both traditional, legacy publishers, and by new content creators. At any moment, an online publisher—for instance, *nyt.com*—may be competing for a finite amount of consumer attention not just with other traditional publishing outlets but with a myriad of content providers across a vast array of other channels—TikTok, Instagram, YouTube, blogging platforms, Twitter, online games, apps, and so forth—putting downward

²⁸ I will try, below, to distinguish which arguments specifically pertain to behavioral online advertising, rather than online advertising tout court.

pressure on revenues per-publisher. Second, the particular form of advertising that consumer data collection has made possible—behavioral targeting via third-party tracking by data intermediaries—has had two opposite effects on publishers’ economic returns. On the one hand, behavioral targeting has made online ads generally more valuable *at the impression level* (targeting is correlated with higher ad conversion rates, and therefore more profitable for publishers, many marketers claim; see Boerman, Kruikemeier, and Zuiderveen Borgesius 2017). On the other hand, behavioral targeting has diminished publishers’ power to match consumers with advertisers, creating an opposite (downward) pressure on publishers’ revenues. Before the rise of behavioral advertising, a merchant selling golf-related products who wanted to advertise to golf-interested consumers may have allocated advertising budget to a specific subset of outlets that counted such consumers among their readers. Online third-party tracking allows advertising intermediaries to target ads to consumers based on the latter’s preferences, regardless of the website, platform, or channel they may be visiting at any given moment (for instance, a visitor to a TikTok video may have been identified as a golf lover, and may be presented with a golf-related ad). Worse (for high-quality, high-cost publishers), a high-value consumer (a reader of the New York Times, for instance) may be targeted while she is visiting lower-quality sites where it is cheaper to advertise (Srinivasan, 2019). These technological opportunities stretch out the supply of advertising spaces—the set of outlets and channels where merchants can find (and buy ad spaces for) interested consumers—shifting the power to match consumers with advertisers away from publishers and toward third-party data intermediaries. With that shift, the ability to extract surplus from advertising transactions also shifts from publishers to intermediaries²⁹—a second source of downward pressure on publishers’ revenues notwithstanding (or, in fact, precisely because of) the advent of more precise ad targeting techniques.

Under this alternative framing of the advertising market (Frame 2), merchants, too, aggressively compete with each to reach consumers with their ads. Before the rise of behavioral advertising, a merchant selling golf-related products intent on advertising to golf-interested consumers may have allocated its advertising budget to related publishing outlets. On those outlets, it would have competed for advertising space with merchants in the same or related industries. Online tracking allows data intermediaries to target ads to a given consumer across platforms based on her multidimensional preferences: the same consumer may be interested in golf, but also in Italian shoes, vacations to Mexico, and cooking lessons. Hence the golf-related merchant interested in reaching a golf-interested consumer may, at any point in time, be competing for the purchase of ad space with a larger array of merchants bidding to show ads about shoes, vacations, and cooking classes. In this sense, behavioral advertising can increase competition for ad space between advertisers (increasing their bids). Such increased competition is not lessened by the fact (which we observed above) that the advertising inventory has also increased, because there is a finite upper boundary to how many ads a consumer can pay attention to and how many products she can buy.

²⁹ In theoretical work, we have shown how an intermediary in a two-sided advertising market can strategically modulate consumer tracking to increase its profit (Marotta, Wu, Zhang, and Acquisti 2022).

Under Frame 2, the economics agents who benefit from tracking and targeting are the intermediaries. The particular features of this technology have increased competition on both sides of the online advertising market, but have favored a concentration of data, and power, in the middle. Under this market structure, control over data translates to control over profits. The large oligopoly intermediaries may be able to extract more surplus from advertising transactions than the aggressively competing stakeholders on either side.

Both Frame 1 and Frame 2 of Figure 2 are based on plausible theoretical arguments. In fact, Bergemann and Bonatti (2022) highlight how digital platforms can generate both dynamics I have highlighted: surplus creation from matching, and surplus extraction from market power. One frame focuses on micro-level effects: per-impression reduction in search costs. The other frame focuses on macro-level effects: the aggregate impact on merchants' and publishers' revenues of competition through oligopoly intermediaries. To some extent, both the search cost reduction story and the oligopoly intermediaries surplus extraction story may in fact be simultaneously occurring. But does either frame have (more) empirical validation? That is, does current research measure how data-driven advertising differentially affects Figure 2's stakeholders? And even if we were to disregard issues of redistribution of wealth among those stakeholders, does current research cleanly identify increases in overall surplus due to those technologies?

The answer is not yet. The degree of attention empirical scholarly research has paid to the different stakeholders in Figure 2 is uneven, possibly because the entities best positioned to measure the value of online advertising (the intermediaries) may not have incentives to conduct, or sponsor, research that may be critical of that value. (This raises obvious concerns over the risk of corporate capture of research in the field.) Oligopoly intermediaries' record-high profits are evident, although the evidence tends to come more often from industry reports rather than empirical scholarly work. Advertising merchants have received most of the research attention, as a substantial amount of work has examined online advertising effectiveness (Boerman et al. 2017). Measuring returns on online advertising spending is notoriously difficult (Johnson 2022),³⁰ and experimental results have shown that online ad spending does not always produce meaningful effects (Blake, Nosko, and Tadelis, 2015). That noted, work in this area has supported the notion that *behaviorally* targeted ads can increase consumer conversion rates and expenditures (Farahat and Bailey 2012; Tadelis et al. 2023). And yet, their impact on merchants' *aggregate* welfare is probably more nuanced than conversion rates associated with specific ad campaigns can capture. This becomes apparent when we contrast per-impression metrics to general-equilibrium metrics. As all merchants can easily engage in this form of advertising, they may, collectively, wind up in zero-sum prisoner's dilemma dynamics. Individually (at the per-impression level), each advertiser experiences a high conversion rate from behaviorally targeted ads. However, each advertiser may have to engage in behaviorally targeting merely to avoid competitors poaching its consumers. In equilibrium, advertisers may

³⁰ For instance: on the one hand, failing to account for endogeneity and selection bias can vastly overestimate the effect of targeted ads, as conversion rates may hide the fact that targeted ads successfully reached those consumers who were, already, highly likely to purchase the product (Aral 2021). On the other hand, merely tracking online conversion rates may miss the effect that online ads may have on offline purchases.

maintain their respective market share, but spend more for it than if they had spent on (for instance) contextual ads.³¹ Alternatively, rather than generating prisoner's dilemma dynamics, online advertising may benefit all participating merchants by expanding consumer demand and consumer spending (possibly via a reduction in consumers' search costs). There is little causal evidence, however, for or against an *aggregate* demand expansion effect of behaviorally targeted advertising, as opposed to it having a mere *redistribution* effect.³²

Publishers—and the impact online advertising and behavioral advertising in particular have on their revenues—are a distant second in terms of scholarly attention. On theoretical grounds, antipodal dynamics are plausible (Chen and Stallaert 2014): behavioral advertising can increase publishers' revenues because merchants are willing to bid more for ads with a higher likelihood of conversion; behavioral advertising can also *reduce* publishers' revenues by creating hyper-targeted subsets of consumers and shrinking competition across merchants to target those consumers, reducing their bids and ultimately publishers' revenues (Levin and Milgrom 2010). Various experiments have shown that behaviorally targeted ads do increase per-impression revenues for publishers relative to non-behaviorally targeted ones. The amount revenues increase, however, varies across studies: from over 50% in a study by Google (Ravichandran and Korula 2019), to about a third of that (18%) in an independent study (Laub, Miller, and Skiera 2022), to even less in a study using an empirical approach similar to Laub et al. (2022) but drawing data from a single large and arguably sophisticated media company (Marotta, Abhishek, and Acquisti 2019).³³ As in the case of empirical studies of privacy regulation, however, these studies individually offer useful data points but are collectively uninformative about the aggregate effect of behavioral advertising (or regulatory restrictions on it) on publishers. Again, we miss the forest for the trees.

First, these studies compare the revenues of targeted and untargeted ads but do not capture the effect of the rising competition publishers face for visitors' attention from an ever-increasing set of advertising channels (and advertising spaces) made possible by behavioral advertising. Therefore, these studies estimate the marginal revenue-increasing effect of targeting advertising space to visitors who actually reached the publisher's site (per-impression returns: Frame 1), but are mute on the *overall* revenue-decreasing effect of competition and the infinite inventory problem (Frame 2). Behavioral advertising giveth, and behavioral advertising taketh away. And yet, to our knowledge, no study has quantified and compared the two contrasting effects.

³¹ This argument is based on the premise that behavioral targeting does work for most merchants. Because of its black box opacity, which allows rampant ad fraud (Hwang, 2020) and makes attribution challenging, its aggregate effect remains murky. Anecdotal evidence suggests, for instance, that after large and small brands alike curtailed their digital spending, they observed no measurable negative impact on downstream business outcomes (see Fou 2021 and Rowe 2021).

³² This argument focuses on behavioral targeting. If we look at the impact of online advertising tout court, we have some indirect evidence: Todri (2022) finds that adblockers decrease a consumer's online spending by 1.45% on average. And yet even this evidence is agnostic regarding aggregate demand effects: it is not known whether the decrease in digital spending implies an overall decrease in demand or, again, a redistribution from online to offline demand, or an increase in other digital spending not captured by the dataset.

³³ Wang, Jiang, and Yang (2023), mentioned above, found that GDPR compliance for a large publisher led to a modest 5.7% decrease in revenue per click.

Second, studies on the impact of regulations or self-regulatory restrictions on tracking and targeting are similarly uninformative about the *aggregate* impact of those interventions, as they only capture the *local*, redistributive effects of particular interventions (Ding, Wu, and Acquisti 2022). By local, we refer to the fact that even the more far-reaching privacy interventions limit tracking and targeting for only some specific subsets of Internet users; for instance, Apple ATT affects users of iOS devices, while the GDPR applies to EU residents who did not consent to tracking or who are visiting websites that invoke the legitimate business interest clause to dispense with visitor consent altogether, and so forth. Those interventions therefore do not impair the tracking of many other categories of users, who thus remain targetable. Hence, those studies are more likely to capture a budget *reallocation* effect of privacy interventions (that is, advertisers reduce ad spending for affected categories and increase it for unaffected categories). They are not designed to study the *aggregate* effects of broadly encompassing regulations and interventions.³⁴ In short, the current scholarly evidence on publishers' revenues captures a valuable but limited piece of the puzzle. That piece holds as much empirical significance as the anecdotal, correlational evidence, coming from publishers' balance sheets, of continuous declines in revenues associated with the rise of behavioral advertising: the revenues of the largest European publishers stagnated over the past ten years, "while Alphabet (Google) and Meta's revenues increased by more than 500% during the same period" (Armitage, Botton, Dejeu-Castang, and Lemoine 2022, p. 9). Globally, newspaper revenue dropped from \$107 billion in 2000 to roughly \$32 billion in 2022 (based on data from GroupM cited in Angwin 2023).³⁵

A legitimate counterpoint to the above argument is that the decline in traditional publishers' revenues has coincided with an increase in the supply (or, at least, in the number of *suppliers*) of other online content (from bloggers to influencers; from TikTok creators to substack writers). The popularity of this content demonstrates a wealthy consumer demand for it. Leaving aside counterfactual questions (could contextual ads support this new content?), how the emergence of new vectors of content dissemination and new creators has affected consumer welfare is harder to establish, as that emergence also raises the thorny issue of content *quality*. As economists, we tend to sidestep those questions by observing that a consumer's demand for a good demonstrates the utility the consumer (expects to) derive from it. Prudent as that may be, it is also unsatisfactory in an online economy which is explicitly designed to employ choice architecture to nudge individuals to consume ever decreasing sound bites of content, and where more and more content is recycled, manipulated, or misleading, if not outright dangerous and malicious (Swire-Thompson and Lazer, 2020). Uncomfortable as the conversation regarding the quality of the new content dissemination and communication channels behavioral advertising is fostering may be, that seems an important conversation to be had, much like the conversation regarding the hurdles of estimating the value consumer accrue from social media consumption

³⁴ Professor Mayer offers a practical example of this argument: "Advertiser bidding behavior would change in a world without behavioral advertising or where it's a rarity. We don't know what those bids would look like, because advertisers just place behavioral bids now. For example, advertisers might start bidding more often and higher prices for demographically, geographically, or contextually targeted ads" (personal communication with the author).

³⁵ See also <https://www.pewresearch.org/journalism/fact-sheet/newspapers/>.

(Brynjolfsson, Collis, and Eggers 2019) versus its negative effects on subjective wellbeing (Allcott, Braghieri, Eichmeyer, and Gentzkow 2020).

Finally, what do we know about consumers? Surprisingly little. Among the stakeholders represented in Figure 2, consumers have received the least attention in scholarly work. The argument for consumers benefiting from online advertising in general and behavioral advertising in particular is more often posited on intuitive arguments than validated with data. In principle, the benefits consumers receive from online advertising may be direct or indirect. The purported direct benefit of *behavioral* advertising is captured in the advertising executive's words quoted earlier in this section: consumers benefit from being presented ads that are more relevant and more interesting. This is a plausible search cost argument: online ads decrease consumers' search cost and present them with offers closer to their preferences, thereby increasing utility. This argument has empirical support: as noted, behaviorally targeted ads are more likely to generate conversions. This argument is also limited, however, and ultimately inconclusive. Search costs are but one factor in consumer utility. Other factors that affect consumer utility from purchasing products advertised to them online include the prices consumers end up paying, the quality of the product they end up buying, the quality of the merchant they end up interacting with, and so forth. Absent counterfactual evidence on the differential effects, along those possible factors, of targeted ads-linked purchases relative to other purchases, it is impossible to draw evidence-based conclusions about the direct consumer welfare effect of behavioral advertising. Only recently has some of that counterfactual evidence started emerging. In a recent working paper, we found that purchasing products from targeted ads, rather than from search results, increased the likelihood of purchasing from a lower quality merchant and increased the expected price of the product (Mustri, Adjerid, and Acquisti 2022). This evidence suggests a potential welfare-decreasing effect of behavioral advertising due to prices and product quality that may countervail the welfare-increasing effect of search cost reduction.

Free access to content and services is often presented as a key *indirect* benefit of the online advertising economy to consumers. To scrutinize the robustness of evidence supporting this claim, it is useful to distinguish between the role of online ads in general and the role of behaviorally targeted ads in particular. The role of online ads in supporting the provision of content and services seems indisputable. Many online services *are* supported via ads. Consumers seem comfortable "paying" for online services with their eyeballs rather than with cash (although a substantial amount of consumers now prefers to block ads altogether³⁶). The role of *behaviorally* targeted ads specifically in the provision of free services and content—and thus the role of consumer tracking and consumer data—is harder to tease out on causal rather than mere correlational grounds, due to the double-edged effect that behavioral advertising can have on the revenues of content creators, which we noted above. In attempting to tease out these effects, extant research leaves us with more questions than answers. Virtually all of today's typologies of online free services and free content already existed on the Internet before the rise of behavioral advertising in (roughly) the mid-2000s. At the time, those services and

³⁶ Interestingly, ad blocker adoption can have *positive* effects on the quantity and variety of articles users consume. See Yan, Miller, and Skiera (2022).

content were supported by contextual or untargeted advertising. To what extent has the dramatic increase in consumer data collection—including the growing ability to identify consumers and link their behaviors across different online and offline contexts—fueled an increase in the provision or quality of free content and services, and to what extent has it fueled an increase in the profit of the matchmakers—the data intermediaries?³⁷ In fact, to what extent is the degradation of privacy an unavoidable price to pay for more or better content, or in fact a necessary condition for innovation?³⁸

Conceptually, these questions amount to a simple economic comparison between the marginal cost of privacy loss and the marginal benefit of data collected. Empirically, answering those questions is anything but simple. We face an array of disparate pieces of anecdotal evidence but lack causal analysis. Anecdotally, the business model of a large number of content or service providers, from online publishers to app developers, does rely on monetizing consumer data. At the same time, a large number of content providers today use hybrid (freemium) models—including online publishers that have been switching to subscription models in both the US and the EU (Lefrere et al. 2022)—perhaps signaling that an insufficient amount of economic value generated from consumer data reaches downstream creators (with the rest, perhaps, being appropriated by data intermediaries). The limited academic research evidence available has produced mixed results. The GDPR may have reduced EU app developers' incentives to create new apps (Janssen, Kesler, Kummer, and Waldfogel 2022), YouTube's removal of personalization for child-directed content following its settlement with the Federal Trade Commission over violations of the Children's Online Privacy Protection Act (COPPA) may have caused child-directed content creators to produce less content (Johnson, Lin, Cooper, and Zhong 2023), and Google's 2019 ban of targeted advertising in Android children's games may have reduced the release of feature updates (Kircher and Foerderer 2023). On the other hand, Apple's introduction of ATT does not appear to have negatively affected the supply of new apps for iOS users (Cheyre et al. 2022), and may have had only a short-term effect on developers' app-monetization strategies (Kesler 2022). Furthermore, the GDPR does not appear to have negatively affected the quantity and quality of EU news and media websites' content (Lefrere et al. 2022).³⁹

The issue considered in this section is not whether economic value can be created from data. That much is clear. The issue is how much we (scholars, regulators, the public) actually and conclusively know about how that value is allocated, and to what extent the claims that new content, services, and even innovation depend on unrestrained data collection (and are damaged by privacy measures) have empirical validation. The analysis presented here

³⁷ For instance, the number of average ads *per video* has seemingly kept increasing on YouTube over time (Berman 2022); to what degree has that increase led to more or better YouTube videos or services?

³⁸ Or, in fact, a sufficient condition? Over the past two decades, Facebook/Meta has gained access to more consumer data than most other companies in history, making significant financial gains from it. To what degree has this unique degree of accumulation of data and wealth led to societally beneficial innovations? See also Ohm (2012).

³⁹ Note that we are focusing here on the effect on content provision (and benefits allocation) of varying amounts of personal information used in online ads. This is related to yet distinct from the discussion of content providers' reliance on online advertising more broadly (see Shiller, Waldfogel, and Ryan 2018).

suggests that these are unresolved questions. This absence of a definite result may in and of itself give us pause.

3.2 Revisiting Assumptions about the Costs of Protection

The second problem with the scale presented in Figure 1 (and with the economic framing of the debate around privacy) lies in the very notion of a beam counterbalancing the value of data and the value of privacy, casting them as opposed rather than parallel policy goals.

The rash of privacy legislation Posner lamented in 1981 and Varian warned us about in 1996 *did* occur. Even though the US still lacks a comprehensive federal privacy law, since the 1980s and the 1990s a myriad of acts, regulations, and enforcement initiatives materialized in the US at both the federal and state levels. And yet, those regulatory efforts did not seem to produce the damages early contributors to the economics of privacy feared. They did not prevent an unprecedented explosion in consumer data collection, the rise of an (estimated: Atikcan and Chalmers 2019) trillion dollar data economy, the growth of new data-driven products and services, and record profits for several data intermediaries. (They also, one may add, failed to sooth consumers' privacy concerns.) Is there a disproportion between economists' fears about privacy protection and its actual impact? Are privacy and analytics (and the extraction of value from data) inherently antithetical, or could both be simultaneously achieved, at least some times, through a combination of technology and targeted policy intervention?

As we noted in Section 2, empirical economic research *has* provided evidence of negative implications of privacy regulation. That evidence, however, has to be carefully contextualized. First, there is parallel evidence that, under certain conditions, privacy regulation can have a positive effect on economic variables (for instance, increase in technology adoption: Adjerid et al. 2016; or identity theft reduction: Romanosky, Telang, and Acquisti 2011) as well as other non-economic policy goals (such as COVID vaccination; see Buckman 2022). We noted in prior work (Acquisti et al. 2020) how this mixed evidence is consistent with extant economic research on the nuanced impact of regulation on innovation:⁴⁰ the direction of the impact will vary based on how particular interventions are designed, implemented, and enforced (BERR 2008).

Second (and with exceptions, as usual: see, for instance, Janssen et al. 2022), many of the studies showing a negative economic impact of privacy regulation ultimately report effects that are precisely identified but small in magnitude. Even a major regulation such as the GDPR has been shown to have produced a combination of diverse effects (Johnson 2023), including negative but modest (Wang, Jiang, and Yang 2023), and even null. (Several possible explanations exist, including the regulation not being actually enforced; or being enforced and adhered to, but the decrease in data availability not causing the downstream damages some

⁴⁰ In the context of environmental protection, Porter (1991) proposed that strict regulations may incentivize innovations and produce efficiency gains. Shao et al. (2020) review the body of literature that over the years developed around the "Porter hypothesis" and find that the impacts of environmental regulation on innovation behavior are complex and include the creation of new technologies, products, and systems.

economists had predicted: see Lefrere et al. 2022.) The same appears to be happening with Apple ATT (see Section 3.1).

Third (and again with exceptions: consider Miller and Tucker 2009), a sizable portion of the literature in this area has focused on regulations' direct impact on business metrics (for instance, reduction in advertising effectiveness, or reduction in the supply of new apps following the GDPR) and has assumed or extrapolated, but not actually measured, downstream welfare effects on consumers (for instance, a reduction in consumer welfare due to less precisely targeted ads or a reduction in their usage of or satisfaction with available apps).

Fourth, some of the literature has focused on local effects rather than general equilibrium effects. We noted above (Section 3.1) that much of the work on restrictions on behavioral targeting are uninformative about the general impact of those restrictions because they capture the effect of local interventions that will affect some audiences and not others and will therefore allow advertisers to reallocate budgets from one entity to another.

Fifth, much of this literature focuses on short-term effects of regulation, from a few months to a few years. The reasons are various and valid, such as producing timely results and identifying robust causal links. But the result is an emphasis on the short-term impact of regulatory shocks (which includes costs that businesses incur as they adapt to new technological and legal frameworks), rather than comprehensive analyses of long-term effects of different privacy regimes. As we noted in Acquisti et al. (2020), the short-term focus is likely to miss the long-term downstream effects of increased consumer protection and of competition and innovation in privacy between firms.

Sixth, the literature has so far by and large ignored the role of privacy-enhancing technologies (Goldberg 2007) and, in particular, privacy-preserving analytics (PPAs), by which I refer to statistical and cryptographic techniques—from homomorphic encryption to differential privacy (Iezzi 2020)—that make it possible to analyze and extract value from data while, to some degree,⁴¹ protecting privacy. Granted, there is no free lunch: as we noted, both privacy and the lack of privacy are redistributive (the interests over data of different stakeholders are not necessarily *ex ante* aligned), and reducing the granularity of data can be costly, as it can reduce its value. But research suggests that those costs may be minimized by careful interventions (Abowd and Schmutte 2019). In recent work, we considered how the application of differentially private mechanisms to census data affects educational funding calculations (Steed et al. 2022).

⁴¹ Privacy-preserving analytics (and, more broadly, privacy-enhancing technologies) can help to some degree, but are no panacea, because processes such as anonymization or data aggregation can mask individual identities or even protect some types of personal information without necessarily averting downstream privacy harm. Consider Google Topics, a framework for interest-based advertising that does without third-party cookies and cross-device tracking (<https://blog.google/products/chrome/get-know-new-topics-api-privacy-sandbox/>). Professor Cheyre writes: "It can be privacy preserving, but it may not change how targeting ultimately operates in the online advertising ecosystem" (personal communication to the author) - that is, the fact that, even when their identities are nominally protected, individuals may be targeted with offers that may or may not be beneficial to them. Furthermore, doubts have been raised about the extent to which privacy measures (such as Apple ATT or Google Topics) materially enhance or will enhance consumer protection or act as tools for increasing control over a market (Sokol and Zhu 2021). This is a valid concern, but its root cause should not be confused: these dynamics are not inherent to privacy protection *per se*, but to specific measures firms may implement to increase market power under the veil of privacy protection.

We found that funding misallocations due to the use of a differentially private mechanism do occur but are marginal compared to much larger misallocations due to existing data error. In addition, we found that a number of simple policy interventions or reforms could reduce the misallocation due to both privacy mechanisms and data errors. Ultimately, the cost (in terms of funding misallocations) due to privacy interventions may be mitigated with proper policy design. One implication of this research is that, before worrying about the alleged costs of privacy protections, it may be prudent to consider whether other steps (such as reduction in data error and noise) may improve statistical practice.

3.3 Tackling the Aggregation Problem and the Economic Dark Matter

The third and final problem in the economic framing of the debate around privacy consists in the lack of adequate measurements of harms from lost privacy—the right-side pan in Figure 1.

In Section 2, I argued that the economics of privacy has, with few exceptions, bypassed all but a handful of the harm of privacy invasions and the benefits of privacy protection. This creates a knowledge gap that hampers evidence-based policymaking. Worse, by stacking tangible economic benefits of data against intangible, unmeasured benefits of abstract concepts such as autonomy or freedom, the scale (and thus the economic debate around privacy) is vitiated by an inherent asymmetry between salient and measurable metrics contrasted to no less important but less salient, less direct, and less tangible factors. The framing therefore emphasizes the importance of one side over the other.⁴²

The scale presented in Figure 1 (and the economic framing of the privacy debate it reflects) is thus flawed not merely on moral grounds (that is, on account of its failure to consider what as economists we may consider “paternalistic” values, such as the moral foundations for privacy protection). The scale is flawed on *economic* grounds, because it misses the “economic dark matter” (Acquisti et al. 2020): the vast evidence of privacy harm we discussed in Section 2 and exemplified through four scenarios.

Whether it is prudent or advisable to measure that economic dark matter is a valid question. The wisdom of considering certain values untradeable (and, in our context, of approaching privacy as a human right when considering regulation, and accept negative changes in some business metrics—when and if they materialize—as the price to pay for those values) lies, precisely, in the knowledge that those values are essential to the functioning of a society even though they may not be (on first analysis) economically measurable or economically efficient. Policymakers (and, more broadly, the public debate around privacy) are therefore stuck in a seemingly unresolvable dilemma. On the one hand, they are expected to calculate the *net* harm of privacy invasions before a market failure is deemed sufficiently alarming to justify policy intervention (Jin and Stivers 2017). On the other hand, economic research is currently failing policymakers,

⁴² The differential privacy community faces a similar problem: “Because of the way [differential privacy mathematics] frames privacy loss through [privacy loss budgets], disclosure risks can appear abstract and difficult to interpret. By contrast, the effects of setting a [privacy loss budget] on downstream data utility are more easily tracked. This asymmetry can privilege data utility as the driving force behind how [privacy loss budgets] are allocated to different queries. We refer to this problem in this section as “the allocation dilemma” (Seeman and Susser 2022).

because, by sidestepping privacy harm and not properly scrutinizing the allocation of benefits from data, it is not measuring net harms. But if quantifying net harm is hard or even formidably difficult, what should guide policymakers? I suggest some possible strategies in the next section.

3.4 Changing the Frame of the Privacy Economic Debate

So far, in Section 3, I have used an economic perspective to highlight systemic problems with the current economic framing of the privacy debate. I have remarked on the paucity of evidence on the allocation of benefits from data; I have emphasized the lack of adequate research on the economic harm of privacy loss; and I have questioned the very premise of construing the debate as a contest between value of privacy and value of data. In short, I have questioned the scientific grounding for the framing. Conversely, I have presented other evidence: consumers care for privacy and act to protect it; yet, economic and behavioral hurdles make it infeasible for individuals to adequately manage their privacy in the online marketplace; the costs of regulatory corrections to those hurdles may be overblown in the current debate; in fact, economic research has bypassed a massive amount of privacy costs, and the evidence that current equilibria ensure fair allocation of benefits from data is scant; furthermore, tools are available to allow both data analytics and privacy protection.

If this critique has merit, it may suggest a way forward in the economic debate around privacy that alters its framing and changes the burden of proof of the arguments around it. Rather than uncritically accepting the current way of framing the debate (*Privacy protection is often costly and at worst inefficient; unless one can demonstrate quantifiable privacy harms, what need is there for government intervention and regulation?*), we could ask instead: *What is the evidence that current products and services cannot be provided in more privacy-preserving manners, and that new privacy-preserving systems and processes cannot efficiently replace current ones?* This is, in essence, a call for turning the tables in the economic debate around privacy. To reach that lofty goal, we need to foster those nascent lines of enquiry I have cited throughout the manuscript—those that tackle new, difficult, and less-studied research questions around the complex interplays of privacy and economic value.

We need to better understand the harms of privacy loss: How do we help consumers and policymakers process the current asymmetry between tangible benefits of data and intangible harms of privacy? Can we (and should we) calculate the economic dark matter? If so, how do we tackle the “aggregation” problem of privacy harm?

We need to better understand the relationship between data protection and value extraction: What are the downstream (long-term, less obvious), and non-easily quantifiable effects of privacy regulation? What are its beneficial effects? What are the economic effects of the

deployment of privacy-enhancing technologies and privacy-preserving analytics, and how are they distributed to different stakeholders—firms, consumers, society as a whole?

And, ultimately, we need to understand better the allocation of value from data: How is the value of data allocated? Who truly benefits from the data economy?

Section 4: Conclusions

The debate we considered in this manuscript is not new. It started over forty years ago. As Posner (1978) decided to sidestep the “definitional problem” and restrict his analysis of privacy to the withholding or concealment of information, Hirshleifer (1980) responded that such a narrow lens of analysis perhaps explained “why our pioneers' attitude toward privacy is—occasional qualifications aside—on the whole hostile. Their tone suggests that we have more privacy than ever before—probably more than is actually good for us or, at any rate, good for economic efficiency and, furthermore, that any person displaying a special desire for privacy is probably just out to hoodwink the rest of us” (p. 650). And while Hirshleifer argued that “the mainland of ‘privacy’ is not the idea of secrecy [...] what we mean by ‘privacy’ is, rather, a concept that might be described as autonomy within society” (p. 649), Posner (1981) rebuffed that “[t]o affix the term privacy to human freedom and autonomy [...] is simply to relabel an old subject—not to identify a new area for economic research” (p. 405)

The rigorous but narrow Posnerian approach to the economic analysis of privacy proved distinctly successful in terms of scholarly research and impact on public discourse. But that very narrow approach, and that success, have laid the foundations for a crisis now emerging on the horizon. The economics of privacy has become more relevant in the debate around privacy, while sidestepping the evidence of significant and far-reaching harms and systemic behavioral hurdles imperiling market solutions to privacy problems. It has bypassed critical research questions outside of a narrow set that has received outsized attention. In doing so, I have argued, the economics of privacy ultimately risks crowding out critical dimensions of privacy not merely from its own field of research, but from the debate over privacy at large, brushing aside non-economic considerations.

That concern, too, is not novel. Hishleifer (1980)'s words appear, today, prophetic:

Recently a new territory has been discovered by economists, the intellectual continent we call ‘privacy.’ The pioneers are our peerless leaders Posner and Stigler whose golden findings have already dazzled the world. It is high time for rattlers and desperadoes—that's the rest of us—to put in an appearance. Of course, I ought to add parenthetically, “new” is relative to one's point of view. Our pioneering economists, like explorers in other places and other times, found aborigines already inhabiting the territory—in this case intellectual primitives, Supreme Court justices and such. Quite properly, our explorers have brushed the natives aside, and I shall follow in that

honorable tradition [...] The first issue I shall address is whether our pioneers have correctly mapped the major features of the "privacy" continent. Have they possibly mistaken a peninsula for the mainland, foothills for a grand sierra, or perhaps even misread their compass so as to reverse north and south? Well, not quite so bad as the last, but I will be contending that the mainland of "privacy" is not the idea of secrecy as our pioneers appear to believe—secrecy is only an outlying peninsula.

Posner won the round, insofar as the economics of privacy adopted a decidedly Posnerian viewpoint. But (to paraphrase the title of a manuscript we cited earlier in this manuscript), at what price? Considering the centrality that information flows have commandeered in our lives and societies over the last four decades, and the extraordinarily far-reaching implications of the control over data and digital boundaries today, the intellectual continent of privacy has become possibly even vaster than Hirshleifer himself may have imagined in 1980. And so when we, as economists, narrow our lens of analysis without correspondingly narrowing the scope of our claims, what dramatic shifts in our societies' economic and social imbalances may we be neglecting? Can we do both—maintain the methodological rigor of our research toolkit, but also expand its narrow horizon of investigation? Will we be able to alter the framing of our research (and the debate around privacy) by accounting for the rich privacy theorizing from other social sciences, and by admitting that a drive for privacy is not inherently antithetical to the extraction of societal benefits from data, since we have technologies and strategies to often allow one and the other?

Posner (1981) wrote that "here as in other areas of nonmarket behavior the economist has a distinctive and valuable contribution to make to social science scholarship" (p. 408). We agree. Used as a complement to the scholarship of other disciplines, the economics of privacy has much to contribute. Used with hubris, mistaking the outlying peninsula for the continent, the economics of privacy risks success at the expense of impoverishing the public debate over privacy; or risks demise by rendering itself decreasingly relevant to it. There is another way, which consists in focusing on a different set of research questions that brave new pioneers in the field may dare to explore, and challenging the way we frame this debate. The economics of privacy is at a crossroads.

References

- Abowd, J.M. and Schmutte, I.M., 2019. An economic analysis of privacy protection and statistical accuracy as social choices. *American Economic Review*, 109(1), pp.171-202.
- Acquisti, A., Adjerid, I., Balebako, R., Brandimarte, L., Cranor, L.F., Komanduri, S., Leon, P.G., Sadeh, N., Schaub, F., Sleeper, M. and Wang, Y., 2017. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)*, 50(3), pp.1-41.
- Acquisti, A., Brandimarte, L. and Hancock, J., 2022. How privacy's past may shape its future. *Science*, 375(6578), pp.270-272.
- Acquisti, A., Brandimarte, L. and Loewenstein, G., 2015. Privacy and human behavior in the age of information. *Science*, 347(6221), pp.509-514.

- Acquisti, A., Brandimarte, L. and Loewenstein, G., 2020. Secrets and likes: The drive for privacy and the difficulty of achieving it in the digital age. *Journal of Consumer Psychology*, 30(4), pp.736-758.
- Acquisti, A., John, L.K. and Loewenstein, G., 2013. What is privacy worth?. *The Journal of Legal Studies*, 42(2), pp.249-274.
- Acquisti, A., Taylor, C. and Wagman, L., 2016. The economics of privacy. *Journal of economic Literature*, 54(2), pp.442-92.
- Acquisti, A. and Varian, H.R., 2005. Conditioning prices on purchase history. *Marketing Science*, 24(3), pp.367-381.
- Adjerid, I., Acquisti, A., Telang, R., Padman, R. and Adler-Milstein, J., 2016. The impact of privacy regulation and technology incentives: The case of health information exchanges. *Management Science*, 62(4), pp.1042-1063.
- Allcott, H., Braghieri, L., Eichmeyer, S. and Gentzkow, M., 2020. The welfare effects of social media. *American Economic Review*, 110(3), pp.629-676.
- Atikcan, E.Ö. and Chalmers, A.W., 2019. Choosing lobbying sides: The general data protection regulation of the European Union. *Journal of Public Policy*, 39(4), pp.543-564.
- Altman, I., 1975. *The environment and social behavior: privacy, personal space, territory, and crowding*. Brooks/Cole Publishing Company.
- Altman, I., 1976. Privacy: A conceptual analysis. *Environment and behavior*, 8(1), pp.7-29.
- Altman, I., 1977. Privacy regulation: Culturally universal or culturally specific?. *Journal of social issues*, 33(3), pp.66-84.
- Angwin, J., 2023. If It's Advertised to You Online, You Probably Shouldn't Buy It. Here's Why. *New York Times*, April 6, 2023, <https://www.nytimes.com/2023/04/06/opinion/online-advertising-privacy-data-surveillance-consumer-quality.html>.
- Aral, S., 2021. What digital advertising gets wrong. *Harvard Business Review*, 19.
- Aral, S. and Eckles, D., 2019. Protecting elections from social media manipulation. *Science*, 365(6456), pp.858-861.
- Armitage, C., Botton, N., Dejeu-Castang, L., and Lemoine, L. 2022. Study on the impact of recent developments in digital advertising on privacy, publishers and advertisers. *European Commission Final Report*.
- Arrieta-Ibarra, I., Goff, L., Jiménez-Hernández, D., Lanier, J. and Weyl, E.G., 2018, May. Should we treat data as labor? Moving beyond "free". In *AEA Papers and Proceedings* (Vol. 108, pp. 38-42).
- Athey, S., Catalini, C. and Tucker, C., 2017. *The digital privacy paradox: Small money, small costs, small talk* (No. w23488). National Bureau of Economic Research.
- Atikcan, E.Ö. and Chalmers, A.W., 2019. Choosing lobbying sides: The general data protection regulation of the European Union. *Journal of Public Policy*, 39(4), pp.543-564.
- Baker, C.E., 1977. Posner's Privacy Mystery and the Failure of Economic Analysis of Law. *Ga. L. Rev.*, 12, p.475.

- Bao, T., Liang, B. and Riyanto, Y.E., 2021. Unpacking the negative welfare effect of social media: Evidence from a large scale nationally representative time-use survey in China. *China Economic Review*, 69, p.101650.
- Becker, G.S., 1980. Privacy and malfeasance: A comment. *The Journal of Legal Studies*, 9(4), pp.823-826.
- Bergemann, D. and Bonatti, A., 2022. Data, Competition, and Digital Platforms. *Working paper*.
- Berman, M. 2022. Why Does YouTube Have So Many Ads in 2022? *Programming Insider*, April 6, 2022, <https://programminginsider.com/why-does-youtube-have-so-many-ads-in-2022/>.
- BERR (Department for Business, Enterprise, and Regulatory Reform), 2008. Regulation and innovation: Evidence and policy implications. BERR Economics Paper no. 4. BERR, UK.
- Blake, T., Nosko, C. and Tadelis, S., 2015. Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1), pp.155-174.
- Bleier, A., Goldfarb, A. and Tucker, C., 2020. Consumer privacy and the future of data-based innovation and marketing. *International Journal of Research in Marketing*, 37(3), pp.466-480.
- Bloustein, E.J., 1977. Privacy is dear at any price: A response to Professor Posner's economic theory. *Ga. L. Rev.*, 12, p.429.
- Boerman, S.C., Kruijkemeier, S. and Zuiderveen Borgesius, F.J., 2017. Online behavioral advertising: A literature review and research agenda. *Journal of advertising*, 46(3), pp.363-376.
- Borgolte, K. and Feamster, N., 2020, April. Understanding the performance costs and benefits of privacy-focused browser extensions. In *Proceedings of The Web Conference 2020* (pp. 2275-2286).
- Bradshaw, S. and Howard, P.N., 2018. Challenging truth and trust: A global inventory of organized social media manipulation. *The computational propaganda project*, 1, pp.1-26.
- Brynjolfsson, E., Collis, A. and Eggers, F., 2019. Using massive online choice experiments to measure changes in well-being. *Proceedings of the National Academy of Sciences*, 116(15), pp.7250-7255.
- Buckman, J.R., Adjerid, I. and Tucker, C., 2022. Privacy Regulation and Barriers to Public Health. *Management Science*.
- Bruns, A., 2021. Echo chambers? Filter bubbles? The misleading metaphors that obscure the real problem. In *Hate speech and polarization in participatory society* (pp. 33-48). Routledge.
- Calo, R., 2011. The boundaries of privacy harm. *Ind. LJ*, 86, p.1131.
- Cecere, G., Le Guel, F., Manant, M. and Soulié, N., 2017. *The economics of privacy*. HAL.
- Chen, J. and Stallaert, J., 2014. An economic analysis of online advertising using behavioral targeting. *Mis Quarterly*, 38(2), pp.429-A7.
- Cheyre, C., Leyden, B., Baviskar, S., and Acquisti, A., 2022. The Impact of Apple Tracking Transparency Framework on the App Ecosystem. *Working Paper presented at WISE*.
- Christou, G. and Rashid, I., 2021. Interest group lobbying in the European Union: privacy, data protection and the right to be forgotten. *Comparative European Politics*, 19(3), pp.380-400.

- Citron, D.K. and Solove, D.J., 2022. Privacy harms. *BUL Rev.*, 102, p.793.
- Colnago, J., Cranor, L.F. and Acquisti, A., 2023. Is There a Reverse Privacy Paradox? An Exploratory Analysis of Gaps Between Privacy Perspectives and Privacy-Seeking Behaviors. *Proceedings on Privacy Enhancing Technologies*, 1, pp.455-476.
- Cooper, J. 2023. Does Privacy Want To Unravel? *Harvard Journal of Law & Technology*, forthcoming.
- Derksen, L., McGahan, A. and Pongeluppe, L., 2022. Privacy at What Cost? Using Electronic Medical Records to Recover Lapsed Patients Into HIV Care. *NBER Workshop on the Economics of Privacy*
- Dienlin, T. and Trepte, S., 2015. Is the privacy paradox a relic of the past? An in-depth analysis of privacy attitudes and privacy behaviors. *European journal of social psychology*, 45(3), pp.285-297.
- Ding, Z., Wu, Y., and Acquisti, A., 2022. Regulation of Targeted Advertising: Profit Implications for Ad Intermediaries and Publishers. *Working Paper presented at WISE*.
- Farahat, A. and Bailey, M.C., 2012, April. How effective is targeted advertising?. In *Proceedings of the 21st international conference on World Wide Web* (pp. 111-120).
- Farrell, J., 2012. Can privacy be just another good. *J. on Telecomm. & High Tech. L.*, 10, p.251.
- Fou, A., 2021. When Big Brands Stopped Spending On Digital Ads, Nothing Happened. Why? *Forbes*, Jan 2, 2021, <https://www.forbes.com/sites/augustinefou/2021/01/02/when-big-brands-stopped-spending-on-digital-ads-nothing-happened-why/?sh=14736b151166>.
- Goldberg, I., 2007. Privacy-Enhancing Technologies for the Internet III: Ten Years Later. *Digital Privacy: Theory, Technologies, and Practices*, p.1.
- Goldberg, S., Johnson, G. and Shriver, S., 2023. Regulating Privacy Online: An Economic Evaluation of the GDPR. *American Economic Journal: Economic Policy*, forthcoming.
- Goldfarb, A. and Que, V.F., 2023. The Economics of Digital Privacy. *Annual Review of Economics*, 15.
- Goldfarb, A. and Tucker, C.E., 2011. Privacy regulation and online advertising. *Management science*, 57(1), pp.57-71.
- Grossklags, J. and Acquisti, A., 2007, June. When 25 Cents is Too Much: An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information. In *WEIS*.
- Guess, A.M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., Gentzkow, M. and González-Bailón, S., 2023. Reshares on social media amplify political news but do not detectably affect beliefs or opinions. *Science*, 381(6656), pp.404-408.
- Hermalin, B.E. and Katz, M.L., 2006. Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quantitative marketing and economics*, 4(3), pp.209-239.
- Hixson, R.F., 1987. *Privacy in a Public Society: Human Rights in Conflict*. New York: Oxford University Press.
- Hirshleifer, J., 1971. The private and social value of information and the reward to inventive activity. *The American Economic Review*, 61(4), pp. 541-556.
- Hirshleifer, J., 1980. Privacy: Its origin, function, and future. *The Journal of Legal Studies*, 9(4), pp.649-664.
- Hui, K.L. and Png, I.P.L., 2006. The Economics of Privacy. In *Economics and Information Systems*.
- Hwang, T., 2020. *Subprime attention crisis: Advertising and the time bomb at the heart of the Internet*.

FSG originals.

- Iezzi, M., 2020, December. Practical privacy-preserving data science with homomorphic encryption: an overview. In *2020 IEEE International Conference on Big Data (Big Data)* (pp. 3979-3988). IEEE.
- Janssen, R., Kesler, R., Kummer, M.E. and Waldfogel, J., 2022. *GDPR and the lost generation of innovative apps* (No. w30028). National Bureau of Economic Research.
- Jia, J., Jin, G.Z. and Wagman, L., 2021. The short-run effects of the general data protection regulation on technology venture investment. *Marketing Science*, 40(4), pp.661-684.
- Jin, G.Z. and Stivers, A., 2017. Protecting consumers in privacy and data security: A perspective of information economics. *Available at SSRN 3006172*.
- Johnson, G.A., 2022. Inferno: A guide to field experiments in online display advertising. *Journal of Economics & Management Strategy*.
- Johnson, G., 2023. Economic research on privacy regulation: Lessons from the GDPR and beyond.
- Johnson, G., Lin, T., Cooper, J.C. and Zhong, L., 2023. COPPAcalypse? The Youtube Settlement's Impact on Kids Content. *The Youtube Settlement's Impact on Kids Content (April 26, 2023)*.
- Jones, C.I. and Tonetti, C., 2020. Nonrivalry and the Economics of Data. *American Economic Review*, 110(9), pp.2819-58.
- Kesler, R., 2022. The Impact of Apple's App Tracking Transparency on App Monetization. *Available at SSRN 4090786*.
- Kircher, T. and Foerderer, J., 2023. Ban Targeted Advertising? An Empirical Investigation of the Consequences for App Development. *Management Science*.
- Laub, R., Miller, K.M. and Skiera, B., 2022. The Economic Value of User Tracking for Publishers. *Available at SSRN 4251233*.
- Laudon, K.C., 1996. Markets and privacy. *Communications of the ACM*, 39(9), pp.92-104.
- Lee, Y.S. and Weber, R., 2021. "Revealed Privacy Preferences: Are Privacy Choices Rational?" *Working paper*.
<https://www.dropbox.com/s/w6q5v5dzpsqferw/Revealed%20Privacy%20Preferences%202021-12-10.pdf?dl=0>
- Lefrere, V., Warberg, L., Cheyre, C., Marotta, V. and Acquisti, A., 2022. Does Privacy Regulation Harm Content Providers? A Longitudinal Analysis of the Impact of the GDPR. *NBER Workshop on the Economics of Privacy*.
- Levin, J. and Milgrom, P., 2010. Online advertising: Heterogeneity and conflation in market design. *American Economic Review*, 100(2), pp.603-07.
- Lin, T., 2022. Valuing intrinsic and instrumental preferences for privacy. *Marketing Science*.
- Liu, Y., Gummadi, K.P., Krishnamurthy, B. and Mislove, A., 2011, Analyzing Facebook privacy settings: user expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference* (pp. 61-70).
- Madden, M., 2012. Privacy management on social media sites. *Pew Internet Report*, 24, pp.1-20.

- Marotta, V., Abhishek, V. and Acquisti, A., 2019. Online tracking and publishers' revenues: An empirical analysis. In *Workshop on the Economics of Information Security*.
- Marotta, V., Wu, Y., Zhang, K. and Acquisti, A., 2022. The welfare impact of targeted advertising technologies. *Information Systems Research*, 33(1), pp.131-151.
- Marthews, A. and Tucker, C.E., 2017. Government surveillance and internet search behavior. *Available at SSRN 2412564*.
- McDonald, A.M. and Cranor, L.F., 2008. The cost of reading privacy policies. *Isjlp*, 4, p.543.
- McDonald, A.M. and Cranor, L.F., 2010. Americans' attitudes about internet behavioral advertising practices. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society* (pp. 63-72).
- Miller, A.R. and Tucker, C.E., 2011. Can health care information technology save babies?. *Journal of Political Economy*, 119(2), pp.289-324.
- Miller, A.R. and Tucker, C., 2018. Privacy protection, personalized medicine, and genetic testing. *Management Science*, 64(10), pp.4648-4668.
- Nissenbaum, H., 2004. Privacy as contextual integrity. *Wash. L. Rev.*, 79, p.119.
- Mustri, E.A.S., Adjerid, I. and Acquisti, A., 2022. Behavioral advertising and consumer welfare: An empirical investigation. *Federal Trade Commission PrivacyCon*.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4398428
- Neumann, N., Tucker, C.E. and Whitfield, T., 2019. Frontiers: How effective is third-party consumer profiling? Evidence from field studies. *Marketing Science*, 38(6), pp.918-926.
- Noam, E.M. and Parker, E.C., 1995. Privacy in Telecommunications: Markets, Rights and Regulations. *New Telecom Quarterly*.
- Norberg, P.A., Horne, D.R. and Horne, D.A., 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of consumer affairs*, 41(1), pp.100-126.
- Nyhan, B., Settle, J., Thorson, E., Wojcieszak, M., Barberá, P., Chen, A.Y., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D. and Freelon, D., 2023. Like-minded sources on Facebook are prevalent but not polarizing. *Nature*, pp.1-8.
- Ohm, P., 2012. The underwhelming benefits of big data. *U. Pa. L. Rev. PENNumbra*, 161, p.339.
- Olson, M.L. 1965. *The Logic of Collective Action*. Cambridge, Mass.: Harvard University Press.
- Pani, L., 2000. Is there an evolutionary mismatch between the normal physiology of the human dopaminergic system and current environmental conditions in industrialized countries?. *Molecular Psychiatry*, 5(5), pp.467-475.
- Porter, M.E., 1991. America's green strategy. *Scientific American*, 264(4), 193-246.
- Posner, R.A., 1977. The right of privacy. *Ga. L. Rev.*, 12, p.393.
- Posner, R.A., 1978. Economic theory of privacy. *Regulation*, 2, p.19.
- Posner, R.A., 1981. The economics of privacy. *The American economic review*, 71(2), pp.405-409.
- Ravichandran, D. and Korula, N., 2019. Effect of disabling third-party cookies on publisher revenue. *Google White Paper*. [Accessed online at https://services.google.com/fh/files/misc/disabling_third-party_cookies_publisher_revenue.pdf on October 4, 2021].

- Rao, A., Schaub, F., Sadeh, N., Acquisti, A. and Kang, R., 2016. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)* (pp. 77-96).
- Reidenberg, J.R., Breaux, T., Cranor, L.F., French, B., Grannis, A., Graves, J.T., Liu, F., McDonald, A., Norton, T.B. and Ramanath, R., 2015. Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Tech. LJ*, 30, p.39.
- Rowe, A., 2021. How Uber's Ad Fraud Lawsuit Highlights a Billion-Dollar Brand Problem. Tech.co, Jan 4, 2021, <https://tech.co/news/uber-ad-fraud-brand-problem>.
- Romanosky, S., Telang, R. and Acquisti, A., 2011. Do data breach disclosure laws reduce identity theft?. *Journal of Policy Analysis and Management*, 30(2), pp.256-286.
- Shao, S., Hu, Z., Cao, J., Yang, L. and Guan, D., 2020. Environmental regulation and enterprise innovation: a review. *Business strategy and the environment*, 29(3), pp.1465-1478.
- Shiller, B., Waldfoegel, J. and Ryan, J., 2018. The effect of ad blocking on website traffic and quality. *The RAND Journal of Economics*, 49(1), pp.43-63.
- Seeman, J. and Susser, D., 2022. Between Privacy and Utility: On Differential Privacy in Theory and Practice. Available at SSRN 4283836.
- Skiera, B., Miller, K. and Jin, Y., 2022. *The impact of the General Data Protection Regulation (GDPR) on the online advertising market*. Bernd Skiera.
- Solove, D.J., 2006. A taxonomy of privacy. *University of Pennsylvania law review*, pp.477-564.
- Solove, D.J., 2007. I've got nothing to hide and other misunderstandings of privacy. *San Diego L. Rev.*, 44, p.745.
- Solove, D.J., 2021. The myth of the privacy paradox. *Geo. Wash. L. Rev.*, 89, p.1.
- Sokol, D.D. and Zhu, F., 2021. Harming competition and consumers under the guise of protecting privacy: An analysis of Apple's iOS 14 policy updates. *Cornell L. Rev. Online*, 107, p.94.
- Srinivasan, D., 2019. The antitrust case against Facebook: A monopolist's journey towards pervasive surveillance in spite of consumers' preference for privacy. *Berkeley Bus. LJ*, 16, p.39.
- Spiekermann, S., Acquisti, A., Böhme, R. and Hui, K.L., 2015. The challenges of personal data markets and privacy. *Electronic markets*, 25(2), pp.161-167.
- Steed, R., Liu, T., Wu, Z.S. and Acquisti, A., 2022. Policy impacts of statistical uncertainty and privacy. *Science*, 377(6609), pp.928-931.
- Stigler, G.J., 1980. An introduction to privacy in economics and politics. *The Journal of Legal Studies*, 9(4), pp.623-644.
- Stigler, G.J. and Becker, G.S., 1977. De gustibus non est disputandum. *The American Economic Review*, 67(2), pp.76-90.
- Strahilevitz, L.J. and Kugler, M.B., 2016. Is privacy policy language irrelevant to consumers?. *The Journal of Legal Studies*, 45(S2), pp.S69-S95.
- Strahilevitz, L. and Liu, L.Y., 2022. Cash Substitution and Deferred Consumption as Data Breach Harms.

University of Chicago Coase-Sandor Institute for Law & Economics Research Paper, (963).

Stutzman, F.D., Gross, R. and Acquisti, A., 2013. Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of privacy and confidentiality*, 4(2), p.2.

Swire-Thompson, B. and Lazer, D., 2020. Public health and online misinformation: challenges and recommendations. *Annu Rev Public Health*, 41(1), pp.433-451.

Tadelis, S., Hooton, C., Manjeer, U., Deisenroth, D., Wernerfelt, N., Dadson, N. and Greenbaum, L., 2023. *Learning, Sophistication, and the Returns to Advertising: Implications for Differences in Firm Performance* (No. w31201). National Bureau of Economic Research.

Taylor, C.R., 2004. Consumer privacy and the market for customer information. *RAND Journal of Economics*, pp.631-650.

Todri, V., 2022. The impact of ad-blockers on online consumer behavior. *Marketing Science*, 41(1), pp.7-18.

Tomaino, G., Wertebroch, K. and Walters, D.J., 2021. Intransitivity of consumer preferences for privacy. *Working paper*.

Turow, J., Hennessy, M. and Draper, N., 2018. Persistent misperceptions: Americans' misplaced confidence in privacy policies, 2003–2015. *Journal of Broadcasting & Electronic Media*, 62(3), pp.461-478.

Vuorre, M. and Przybylski, A.K., 2023. Estimating the association between Facebook adoption and well-being in 72 countries. *Royal Society Open Science*, 10(8), p.221451.

Varian, H.R., 1996. Economic aspects of personal privacy, privacy and self-regulation in the information age. *National Telecommunications and Information Administration Report*. Reprinted in Lehr, W.H. and Pupillo, L.M. eds., 2009. *Internet policy and economics: Challenges and perspectives*.

Yan, S., Miller, K.M. and Skiera, B., 2022. How does the adoption of ad blockers affect news consumption?. *Journal of Marketing Research*, 59(5), pp.1002-1018.

Wang, P., Jiang, L. and Yang, J., 2023. The Early Impact of GDPR Compliance on Display Advertising: The Case of an Ad Publisher. *Journal of Marketing Research*.

Wernerfelt, N., Tuchman, A., Shapiro, B. and Moakler, R., 2022. Estimating the value of offsite data to advertisers on meta. University of Chicago, Becker Friedman Institute for Economics Working Paper, (114).

Westin, A. 1967. *Privacy and Freedom*. Ig Publishing.

Wickelgren, A.L., 2015. An economic analysis of internet privacy regulation.

Zuboff, S., 2015. Big other: surveillance capitalism and the prospects of an information civilization. *Journal of information technology*, 30(1), pp.75-89.