The Economics of Privacy at a Crossroads

Alessandro Acquisti[1]
Carnegie Mellon University

NBER Tutorial on the Economics of Privacy

Draft, January 18 2023. Comments welcome: acquisti@andrew.cmu.edu

By several accounts, the economics of privacy has grown into a remarkably successful field of research. As the means of collecting and using consumer data have expanded, so has the body of work investigating trade-offs associated with flows of those data. 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.

Barely 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 of the public discourse those other non-economic considerations, 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 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 by considering economic dimensions and research questions that, so far, have received limited attention; and by accounting for the broader scholarship on privacy coming from other disciplines. As a complement rather than a substitute to the contributions of other fields, the economics of privacy may keep thriving, and remain a useful tool for debating 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 current days, and celebrate its successes. In Section 2 I 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 (but not all) current research, which fixates on the economic costs of privacy protection at the expense of resolving a rich array of yet unanswered questions.

[1] Alessandro Acquisti gratefully acknowledges support from the MacArthur Foundation through grant 22-2203-156318-TPI.

Section 1: The Rise of the Economics of Privacy

The economics of privacy is not a novel field of research. It can boast a venerable pedigree. 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). Here, I focus on highlighting a few key milestones. A wave of economic analyses of privacy started appearing near the end of the 1970s and the start of the 1980s. Several of those analyses were produced by Chicago scholars interested in economics and law, such as Posner and Stigler. They were the intellectual "pioneers" (in Hirshleifer (1980)'s wording) who had discovered a "new territory [...] the intellectual continent we call 'privacy'." It is not ungenerous to describe, as Hirshleifer (1980) did, those pioneers' views as "hostile" 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 restriction on data collection - it was about *autonomy within society*. In some sense, the dispute Posner and Hirshleifer commenced four decades ago has never been 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 Lauden started writing again on 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 nascent or impending economic implications of those changes. Varian (1996) diagnosed the link between economics and technology at the root of modern privacy problems: data that was already theoretically public (or at least accessible) in physical format becomes much cheaper to capture, store, and access once digitized, and thus "more" public; as its price lowers, demand for it increases. Noam (1995) wrote about the economic interpretation of encryption. 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 - an idea that has taken different manifestations in the roughly 25 or so years since it was first proposed, and that has been the subject of numerous proposals (from data dividends to data as labor; see Arrieta-Ibarra, Goff, Jiménez-Hernández, Lanier, and Weyl 2018).[2]

[2] In our 2016 review of the economics of privacy in the *Journal of Economic Literature* (Acquisti et al 2016), as well as in recent pieces (Spiekermann, Acquisti, Böhme, Hui 2015; Acquisti, Brandimarte, and Loewenstein 2020), we discuss some of the issues and 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 still do not exist, notwithstanding widespread scholarly and commercial interest. Central among those reasons is 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 users control over co-created data, platforms keep legal and economic control of user data, undermining consumers' ability to leverage parallel "data markets" to protect (or merely commercially benefit from trades over) their personal data.

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. Varian noted that individuals may strategically prefer to share some personal data white protecting other data. For instance, the same consumer may want her preferences to be shared with a merchant (so as to get personalized offers), but not her reservation price (so as to avoid first degree price discrimination). Noam observed that in privacy interactions the rights of different parties collide (and so do their preferences and interests). We can infer from this that it would be naive to expect *ex ante* that the (economic) interests different parties hold over sharing or protecting personal data will be aligned. For instance, the seller has a clear interest in knowing the buyer's reservation price for a good; the buyer has an interest in protecting that information.

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 and storage kept falling; the costs of computation kept dropping; 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" research started expanding and diversifying - although they remained mostly tied to the *informational* dimension of privacy. While in the early 2000 much research on this topic focused on data breaches and price discrimination, over time the topics covered multiplied: from the relationship 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 grown. Their backgrounds have become more diverse too - 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 premiere economic outlets, such as the American Economic Review, the Journal of Political Economy, and RAND. There are anecdotal yet meaningful signals of relevance, too: the publication in JEL of a review of the field, and the hosting at the NBER 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 - by which we often refer to undesired ramifications of regulatory interventions in the market. In this section I discuss the unintended consequences of the undeniable 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 (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 reductionist approach to privacy (pioneered by Posner and Stigler and still common today in 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, yet are no less important.

2.1. How Economists Think of Privacy vs How Privacy Scholars Think of Privacy

Posner's 1977, 1978, and 1981 essays on privacy proved remarkably influential over time, and not merely because of their citation count. Notwithstanding the dramatic growth of this field of research, to some extent the economics of privacy never fully outgrew the framing and directions Posner set for it over 40 years ago.

In the first couple of pages of his 1981 article, Posner makes four remarkable points. First, after having acknowledged different interpretations of the term "privacy," Posner identifies one as the most deserving of economic attention: concealment. In fact (and 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 those deficiencies" and to "invoke a 'right of privacy'." If privacy is 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, the 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 will not claim that Posner's construction of privacy as concealment of *negative* information still influences today the economic scholarship around privacy.[3] The notion of privacy as being about something to hide has been repeatedly debunked (see 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 my own), and how they compare to the theories and findings of some prominent privacy scholars outside the economic domain. Table 1 offers a stylized comparison between Posner's (and, I am going to claim, much albeit not all contemporary mainstream economics') views on privacy and seminal privacy scholars' viewpoints. I

[3] Although it still sporadically shows up in the public debate. During an interview, Eric Schmidt (then Google's CEO) famously countered a question about 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/?__source=vty%7Cinsidegoogle%7C&par=vty.

discuss those differences immediately below, and then the repercussions of those differences in Section 2.2. A caveat applies: the table extrapolates and generalizes by focusing on trends and prominent views. The economics of privacy has grown more diverse over time, hence exceptions to those trends and views exist and, across the manuscript, I offer examples of some.

Economics	Privacy scholarship
Privacy as concealment	Privacy as data protection, but also as control, as boundary regulation, as autonomy within society, …
Privacy is about personal information	Privacy is about personal information, but also decisional autonomy, freedom (including bodily freedom), liberty, dignity, ….
Revealed preferences: Consumers are rational, and do not care *that much* about privacy	Revealed preferences: Consumers care, and act to manage privacy, but they are hindered by economic and behavioral hurdles
Generally skeptical about government regulation	Generally sympathetic towards (or at least open to) government regulation

Table 1

The first difference in Table 1 relates to the definition of privacy. Posner proposed that the concealment of information was privacy's most interesting meaning from an economic standpoint. Since then, explicit or implied references to privacy as concealment or protection of information flows have remained common in both analytical economic perspectives on privacy and in empirical economic works. For instance, in an overview of privacy and economics written while the authors served in the Economic Bureau at the Federal Trade Commission, Jin and Stivers (2017) drew a distinction between privacy processes and privacy outcomes, and after acknowledging that consumers "want [...] to have a certain amount of *control* over the flow [of individual information]" (emphasis added), they defined "an individual's privacy outcome" as "the realized *restriction* on the flow and use of information" (p.1, emphasis added), and noted that "[a]n entity has more privacy as the flow and use of information about it is more restricted" (p. 5). As recently as 2022, at the first NBER Workshop on the Economics of Privacy, in a remarkable study of the use of electronic medical records to prevent AIDS deaths by enabling patient tracing, Derksen, McGahan, and Pongeluppe (2022) equated privacy with 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"). Furthermore, as we discuss further below, a large portion of empirical economic studies on privacy are, in fact, studies of *data protection*.

For privacy scholars, privacy may *include* concealment or data protection, but is something broader, more nuanced, and ultimately quite different. Where economists "often think of privacy preference as generated from the need to protect one's private information in market exchanges," writes Lin (2022, p. 665), "[p]hilosophers often see privacy as an intrinsic value, 'an aspect of human dignity.'" Across other social sciences, privacy has been linked to control, autonomy, boundary regulation, and more (Acquisti,

Brandimarte, and Loewenstein 2015; Solove 2006). Even when narrowly applied to information, "control" is construed as more than protection - it implies the ability to both protect and to share about oneself (Westin 1967). In 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 Altman's 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 the sharing and the protection of 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. As economists, we may feel queasy about studying concepts seemingly as intangible as the regulation of boundaries. And yet, the distinction is highly consequential in terms of how consumer behavior around privacy is interpreted and how regulatory interventions in privacy are perceived by economists. Through the lens of privacy scholarship, for instance, HIV patients being alarmed about medical tracing and rejecting electronic medical records (see example 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 will 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. Contemporaries of Posner had already detected the limiting constraints of his definition, and ascribed to those bounds the "hostility" economics appeared to breed in regards to privacy (among them, Hirshleifer 1980). Some degree of privacy aversion (and in particular aversion to regulatory interventions; see Section 2.2) still permeates today the economic debate. It may be, in part, due to misconstruing privacy as mere data removal, rather than embracing its function in permitting individuals to manage how they connect with the rest of society. By missing this nuance, 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.[4]

The second difference in Table 1 is about the scope of research. Consistent with Posner's focus on privacy as concealment of *information*, the vast majority of economic scholarship in this field has concentrated on the study of *data* protection. While the specific areas of application 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. Similarly, with some exceptions (for instance, Marthews and Tucker 2017), the empirical literature has focused on the economic ramifications of curtailing access to those data through regulation, self-regulation, or technology across diverse application scenarios (see a review in Acquisti et al 2016).

[4] With some exceptions (for instance, Laudon 1996), it is telling - and alarming - that references to the writings of some of the most influential scholars and theorists of privacy, such as Westin and especially Altman, are rare or nearly entirely absent from economic writings.

Privacy economic research has good reasons to focus on information and data. Information assets are those that have become central to the economic calculus of people and organizations, and the novel privacy concerns that have arisen in the recent decades are, at least on first analysis, informational concerns. However, Altman's theory of boundary regulation does not merely apply to informational boundaries, and privacy scholars do not identify privacy with data. Different boundaries between the self and others exist, including spatial, bodily, and decisional. In fact, Altman's boundary regulation theory is expansive enough to connect the vast array and diversity of privacy definitions and dimensions (Solove 2006) that have been proposed in the literature. The practical manifestations or embodiments of those boundaries can take many forms depending on context; what they have in common is the alternating of the opening and closing of the self to others. Which 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, if not entirely ignored, to a great extent sidestepped in the economic debate around privacy.

The third difference in Table 1 captures the divide between economists' conceptualization of privacy behavior or decision-making and that of behavioral privacy scholars. Posner and Stigler looked at consumers' disclosure decisions as economically rational processes, where individuals signal positive traits but want privacy for negative ones. The belief that consumers can make *economically* rational privacy decisions is still reflected nowadays in the interest some economists have demonstrated towards data markets or towards privacy policy making that favors informational interventions to help consumers better navigate privacy trade-offs in the market. It 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). As economists, we are trained to rely on the concept of revealed preferences. If privacy is narrowly construed as *protection* of personal information, and if privacy behavior is economically rational, 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. That evidence would then be interpreted as proof that individuals do not care for privacy *that much*. Results from experiments whose participants exhibit willingness to depart with personal information in exchange for tiny rewards (Athey, Catalini, and Tucker 2017; Grossklags and Acquisti 2007) may be also interpreted to support these conclusions.

Behavioral privacy research, however, presents evidence in contrast with a Posnerian interpretation of purely strategic privacy decision making. First, an extensive body of work has uncovered a large number of hurdles - asymmetric information, 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). Second, behavioral research has provided clear evidence of extensive privacy seeking behavior - online and offline. Writing two decades before the Internet, Altman (1975) had noted that privacy regulating behaviors are ubiquitous and common, and often occur with little conscious awareness. Those behaviors may be invisible to us economists merely because they escape our definitions of privacy. Ordinary examples from our daily lives abound offline and online (see Acquisti et al 2020, from which these examples are taken): we lower our voice or change topic when, as we are engaged in an intimate conversation with someone, a third party approaches; we step aside from a group of friends when we get the call from the doctor's office with

the results of a test; 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. Numerous studies (including self-report surveys, observational field works, and online experiments) complement the anecdotal observations. Here are just a few examples: 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;[5] and a substantial proportion of Internet users worldwide use tools to block unwanted ads popping up on their browsers adblockers (the proportion varies from study to study, from 27% to close to 50%).[6] In fact, in 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), we found that the 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 a majority of randomly picked list of privacy behaviors. We also found that even a majority of those participants who had claimed privacy not to be particularly important to them had, in fact, engaged in those privacy-protective behaviors (Colnago, Cranor, and Acquisti 2023). The empirical behavioral evidence thus suggests that, contra the notion of digital denizens doing little to protect their privacy, contemporary individuals 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), even though they do not *protect* their data *every time* (Acquisti et al 2020). Of course, they don't: privacy, from an Altmanian perspective, is about individuals 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, 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 one's physical space, and to recognize (and react accordingly) friend from stranger or foe, 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. Hence 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 towards privacy.

The fourth difference in Table 1 is about 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 been often skeptical or outright averse to privacy regulation (again,

[5] See https://www.macrumors.com/2021/05/07/most-iphone-users-app-tracking-opt-out/.
[6] See https://www.insiderintelligence.com/content/ad-blocking-growth-is-slowing-down-but-not-going-away.

exceptions exist: see, for instance, Becker 1980). At the very onset 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 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." Roughly another 20 years later, in an exceptionally balanced piece, Jin and Stivers (2017) considered a number of 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 [...] this should help bolster a *healthy market* to facilitate and honor their choice of privacy" (emphasis added to the reference to policy interventions such as informational or educational campaigns that assist market solutions), 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 policymaker would have such 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, once economic damage has been established (emphasis added, as I will go back to the concept of net harm later in this manuscript). These analytical concerns are reflected in the empirical literature. Mimicking Posner's skepticism towards regulatory interventions, a large share of empirical economic research on privacy has focused on documenting the costs and inefficiencies caused by data protection and privacy regulation, as I further discuss in Section 2.2.

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 cells on the left column of Table 1 are causally linked. If privacy is mainly about concealment, then the abundant evidence of online disclosures is proof of weak individual preferences for privacy; and if rational behaviors in the marketplace accurately capture preferences, then privacy regulation is unnecessary or even deleterious. The chain on the right column of Table 1 follows an alternative logic: if privacy is more than concealment and pertains to more than information, evidence of public disclosures is not proof that individuals do not particularly 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 modern digital marketplace; regulation may thus be needed to allow individuals to manage their privacy in a world of endemic information asymmetry and systemic power imbalances.

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 on quantifiable dimensions of privacy which pioneers such as Posner and Stigler proposed. That approach deviates from much of other social sciences' theorizing on privacy. I want to discuss the unintended consequences of that deviation. Because as economists we mischaracterize privacy (or at the very least sidestep the richness of the multiform dimensions of privacy in the literature outside

economics), we spend more time focusing on the tree (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 other than the study of the impact of data protection; from the evidence of widespread consumer privacy harm; and from the implications of privacy behavioral research.

2.2.1 The Disconnect Between Empirical and Theoretical Privacy Research

A first consequence of the narrow (and to a large extent) skeptical economic view of privacy is the disproportionate attention that empirical works have paid to the costs caused by data protection regulation.

While several exceptions exist (and I offer examples below), the most common focus of empirical research in this field has been the quantification of economic inefficiencies and costs arising from data protection and regulation: from reducing hypothetical purchase intentions in response to limitations on targeting (Goldfarb and Tucker 2011), to decreasing the speed of adoption of electronic medical records and technologies that can save infants' life (Miller and Tucker 2011), to reducing ecommerce spending (Goldberg, Johnson, and Shriver 2019) - just to name a few. Individually, these and many other studies are rigorous. In the aggregate, they reveal a disconnect between the dominant empirical analysis and the theoretical privacy economics literature. The disconnect curtails attention towards other critical research questions.

The theoretical privacy literature has repeatedly highlighted highly nuanced economic effects of both data protection and data 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), data or privacy protection can be either welfare-decreasing and welfare-enhancing, depending on context. The nuanced effects in terms of individual welfare are the easiest to illustrate intuitively: Varian (1996) had already pointed out that *not* sharing personal data could both benefit the consumer (when that data was her reservation price) or 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. 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, privacy enhancing technologies; see Section 3) to avoid 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 aggregate-level welfare argument is less intuitive because it can take multiple forms. Several theoretical pieces (also reviewed in Acquisti et al 2016) show how a lack of privacy can decrease (not just increase) aggregate welfare. They range from Hirshleifer (1971)'s over-investment argument (firms do not internalize consumer privacy costs and invest more than socially optimally on data collection), to Hermalin and Katz (2006)'s *ex ante* vs. *ex post* trade efficiency argument (under which the provision of privacy can create welfare-increasing equilibria that otherwise would be destroyed - even when "there is no 'taste' for privacy *per se*"; p. 209). One illustration of Hermalin and Katz's theoretical argument appears, today, prescient: "[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), which we have cited above. HIV patients may dodge tracing precisely because of their (often justified) fear that medical conditions will not be kept confidential. Strong, credible assurances of privacy protection may induce patients to consent to tracing, improving individual and societal wellbeing. In a timely and insightful manuscript, Buckman, Adjerid, and Tucker (2022) similarly find that privacy protection can increase consumer demand for COVID-19 vaccines.

With important exceptions (for instance, Marthews and Tucker 2017; Neumann, Tucker, and Whitfield 2019; and others), 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 (the careful study by Buckman et al 2022 I just cited, and 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 benign: the empirical literature is open-mindedly testing all sorts of theoretical predictions, and it is finding support for those which highlight the costs of data protection; the costs *are* there. A different explanation is self-selection: for various reasons, we tend to pick questions that aim at finding some costs of regulation - and, of course, we *do* find evidence for those costs, since we focus on short-term metrics most likely to capture those costs. Those reasons may include training and 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 (venture capital investments following the enactment of the General Data Protection Regulation, or GDPR: Jia, Jin, and Wagman 2021; app developers' revenues or app supply following the introduction of Apple ATT: Cheyre, Leyden, Baviskar, and Acquisti 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 (see Section 2.2.2), but because they are much more difficult to quantify and to causally link to the regulation itself. 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 muster the exacting peer-reviewing process. And given that economic journals are not frequently accused of being averse to results exposing the unintended effects of pesky regulators, a marginal costs vs marginal benefits analysis can nudge us towards certain research questions instead of others, and leave unaddressed a plurality of legitimate research questions.

2.2.2 The Economic Paradox of Privacy Harm and the Aggregation Problem

A corollary of the dominant attention of empirical economic research 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 privacy harm.

As noted in Section 2.1, the economics of privacy has predominantly focused on informational issues. Accordingly, the literature has focussed 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 over 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 *data* protection. Because of this, many typologies of consumer privacy harm have been ignored by economic research. In fact, the very existence of consumer concerns over privacy has been sometimes source of explicit *bewilderment* in our field: "why people should want to suppress such facts is mysterious from an economic standpoint" (Posner 1981, referring to publicizing facts that have no possible value to potential transacting partners); "the privacy legislation movement remains a puzzle from the economic standpoint" (Posner 1978); "[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" (Wickelgren 2015). To be fair, theoretical work (e.g., Becker 1980, Hermalin and Katz 2006, and 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 actual consumer harm are lacking: by and large, the economic scholarship has sidestepped the vast array and diversity of harm discussed at length in the legal privacy scholarship (Calo 2011, Citron and Solove 2022).[7]

The economic paradox of privacy harm is that measuring it is hard not because of its rarity, but for the opposite reason: privacy harms are ubiquitous, but diverse in form, heterogeneous in likelihood, and varying in magnitude. 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 psychological stigma to rare but catastrophic physical consequences, and more. Commercial surveillance practices that increase - often without individuals' knowledge and consent - the amount of consumer data collected and shared with third-parties ultimately increase the stochastic risk that any *one* of those myriad potential harms may occur. Therefore, while the individual likelihood of any one 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 as a whole. And yet that expected

[7] Lin (2022) estimated and compared intrinsic and instrumental *preferences* (valuations) for privacy. This comparison is similar to Farrell's intermediate/final good distinction, and is different from the measurement of different typologies of consumer *realized privacy harm* we consider here. Such harm is stochastically realized and unpredictable ex ante for the consumer. 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. 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.

cost remains hard to quantify (for scholars, policymakers, and the consumer herself) because of the aggregation problem.

Consider the following scenarios capturing ramifications of consumer data collection.

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 purposes of online advertising. This happens on the vast majority of websites. This transaction cost is minimal at the individual visit level. Multiplied across multiple visits conducted by an individual over time, and across multiple individuals, the aggregated opportunity costs of time lost to trackers is remarkable. This scenario is an example of a widely common (high likelihood) cost that is minimal at the event-level but significant 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, Latice 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."[8] Following the US Supreme Court's overturn of Roe v. Wade on June 24, 2022 with Dobbs v Jackson, 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).[9] This is an example of an event with very 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.[10] The report called out Facebook as "a useful instrument for those seeking to spread hate" (p. 14): the Myanmar military had used Facebook to systematically engage in propaganda against the Rohingya people. It is difficult to overstate the central role the collection and analysis of personal data by social media platforms plays in these dynamics. Algorithms use personal data to select which information to show to which users in order to boost engagement. They are blind to whether that will mean encouraging a visitor to watch one more video about their favorite football team or riling her up with rage over the purported misdeeds of another group of people. This is an example of a very common phenomenon (algorithmic targeting) contributing to an exceedingly rare event with catastrophic individual and societal consequences (genocidal violence).

Scenario 4: Bradshaw and Howard (2018) found "evidence of formally 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 public opinion domestically," including through the use of disinformation campaigns. As in the Myanmar case, personal data play a central role in these operations, especially with misinformation designed to appeal to specific groups. And yet, while social

[8] See https://www.theverge.com/23185081/abortion-data-privacy-roe-v-wade-dobbs-surveillance-period-tracking.
[9] See https://www.pregnancyjusticeus.org/victory-for-latice-fisher-in-mississippi/.
[10] See https://www.ohchr.org/Documents/HRBodies/HRCouncil/FFM-Myanmar/A_HRC_39_64.pdf.

media *may* sway small but ultimately key portions of voters in very close elections, it may be impossible to demonstrably conclude whether and when an election was won or lost due to how unknowing voters' data was used to target them. Considering the far-reaching ramifications (economic and not) of a nation voting one leader over the others, we are left with a discomforting thought: data-driven online campaigns may have potentially staggering, yet impossible to causally demonstrate, downstream effects on the citizenry. Under this scenario, the probability of attack (election manipulation attempts) is high, but the proof that an attack succeeded (that is, it successfully swayed a significant enough number of voters to affect electoral results) may remain elusive. Which means that scholars (and policymakers) face a dual challenge: the downstream economic ramifications of a success of these attacks may be enormous; and yet, they may be so complex and ramified to be essentially both incalculable and indemonstrable.

To emphasize complexity and heterogeneity, the four selected scenarios vary in likelihood, magnitude, and typology of privacy harm. They are particular examples picked from a much 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. So are the outcomes of those usages: the value of information - and thus of privacy - can often be only determined ex post, based on the context in which information is used. Hence those consequences may include both negative (harms) and positive (benefits) externalities from data collection. It may prove therefore hopeless to attempt to aggregate privacy net harm into a single economic estimate.

And yet, none of those hurdles - 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, other than subjective concerns. Because of its peculiar approach to privacy, the economics of privacy has to a great extent sidestepped the evidence of consumer privacy harm. We measure the tip of the iceberg and remain wholly unfamiliar with its mass underwater.

2.2.3 Lessons 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, following Posner's mold, tends to believe in a rational decision-making process through which consumers reveal their (privacy) preferences in the marketplace. Under that account, if we consider privacy to be the mere concealment of personal information, rampant online sharing and willingness to disclose in exchange for small rewards could be taken as evidence of weak consumer preferences for privacy. The hurdles consumers face in making privacy choices are - at times - acknowledged in this literature, but relegated to sidenotes. Informational interventions are presented as viable strategies to assist privacy-conscious consumers.[11]

Some of the key findings from the behavioral literature paint a different picture and support different implications. Consumers do care, and often engage in privacy-managing and privacy-seeking behaviors

[11] Metaphorically and literally. See, for instance, footnote 38 on p. 23 in Jin and Stivers (2017), who acknowledge some of the hurdles consumers face in the market.

(see Section 2.1). But economic and behavioral hurdles, far from being sidenotes or exceptions, are ubiquitous and central in consumer choice. Those hurdles make consumers' desired degrees of privacy unattainable through market interactions. No amount of informational or educational interventions can remedy those systemic barriers.[12]

The hurdles come in many forms. Some are informational. Data - unlike physical goods - can be non-rival (Jones and Tonetti 2020) and non-exclusive, and is subject to secondary use. Varian (1996) had already observed that widespread secondary use of digital data could give rise to externalities. Individuals rarely know or predict the many possible secondary uses of their data (examples in the literature abound: every week, new ways in which personal information is collected or used are discovered, and users' expectations regarding their privacy are often distant from reality: for instance, see Liu, Gummadi, Krishnamurthy, and Mislove 2011). Hence even a consumer who knowingly engaged in a data transaction with another party will ultimately face externalities she is not able to predict, account for, or control as a rational economic agent: one cannot protect against, or make optimal decisions about, something one does not even know is happening.

Other hurdles are behavioral. A vast array of studies has highlighted troves of cognitive and behavioral factors that can affect privacy decision making (Acquisti et al 2015, 2020). Drawing attention to those factors is very far from suggesting that consumer privacy behavior is erratic or irrational, or that privacy choices are 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, Schaub, Sadeh, Acquisti, Kang 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.

Other hurdles are economic in nature. Economic barriers make it overly costly for consumers to comprehensively manage their digital privacy and often, they render privacy options entirely inaccessible. They include learning costs, adoption costs, switching costs, privacy externalities, lack of viable market alternatives; 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. In contrast, mainstream economics' distrust of regulation and confidence in markets' ability to generate efficient equilibria betray its faith in a hyper-rational theory of choice, and ultimately its dismissal of the key implications of behavioral research.

Wary of paternalistic interventions, economists who have sensed the problem faced by consumers have described informational, educational, and transparency interventions as alternative policymaking tools for digital privacy (Jin and Stivers 2017). Yet the behavioral literature also suggests that no amount of education or transparency interventions would help. Notice and consent regimes do not even resolve the

[12] There are several empirical reasons to believe that transparency and consent are at best necessary, but not sufficient, means of privacy management (see related discussion in Acquisti et al 2020 and Acquisti et al 2022).

basic problem of information asymmetry: they are exorbitantly costly for end-users (McDonald and Cranor 2008), unhelpfully ambiguous and therefore unactionable (Reidenberg, Breaux, Cranor, French, Grannis, Graves, Liu, McDonald, Norton, and Ramanath 2015), and crash under the weight of both the myriad privacy notifications, options, and requests consumers are inundated with daily (Almuhimedi, Schaub, Sadeh, Adjerid, Acquisti, Gluck, Cranor, and Agarwal 2015), and our innate bounded rationality. 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 instinctual and natural. Online, it is 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 so-called dark patterns - the design of user interfaces to subtly take advantage of cognitive and behavioral biases and nudge consumers towards more engagement and self-disclosures (Acquisti, Adjerid, Balebako, Brandimarte, Cranor, Komanduri, Leon, Sadeh, Schaub, Sleeper, and Wang 2017).

Related to this discussion is a specific and contentious stream of privacy behavioral work, that has been the object of particular misapprehensions and thus confusion about the implications of behavioral findings: 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. This leads to disagreements over the *paradoxical* (or not) nature 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.[13] A second and more consequential source of confusion is the implicit

[13] 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 contrast over this point is more driven by grammar than by actual empirical comparison of mental states and behaviors. 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 lenses 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 dichotomy through the lenses 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."

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 binary terms: yes or no. This, of course, is folly: answering in general binary terms would require believing that attitudes *always* match behaviors, or never do. Whereas, everything about privacy (including decision making) is 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 onto which people assign the most diverse interpretations based on their own assumptions, and thus the most diverse policy conclusions. One conclusion (which I disagree with) is that the privacy paradox literature demonstrates that people do not really care about privacy, or they do not really know what they want, and therefore no public intervention is needed, other than perhaps some informational intervention. A different conclusion (which I agree with) is that the existence of a gap between mental states and market choices reflects precisely those economic and behavioral hurdles that we have espoused in this section, and that justify 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 an historical 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, which includes the policies that, at any given time, a society accepts as legitimate options.[14] 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 constant of human cultures across history and geography. The same cannot be said of the notion of privacy as a fundamental human right. Construing privacy as a right is a modern development, progressively and unevenly arisen across different cultures over time (Hixson 1987). The notion of privacy as a fundamental right eventually reached sufficient legitimacy to be ingrained in the principles of an economic organization such as the OECD. In its *Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, adopted in 1980, the OECD remarked that privacy protection laws

[14] See https://www.mackinac.org/7504.

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 unauthorized disclosure of such data" (emphasis added).[15] 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."

Over the past 40 years, the rise of the economics of privacy did not merely provide a useful analytical complement to value-grounded views of privacy. By framing data (and privacy) as tradable assets, it may have diminished the currency of notions of privacy as a right. When Posner outright dismissed attempts to link privacy to broader values such as freedom and autonomy ("[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," he wrote in 1981, p. 405; similar concepts are found in his 1977 and 1978 pieces), 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." Yet, Posner's framework flourished within economics, and over time influenced public policy. When the OECD revisited, in 2013, its 1980 Guidelines, the term "fundamental value" had replaced the original "fundamental right."[16] 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, although it was explicitly used in reference to *other* rights, such as freedom of speech, freedom of the 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 assessments. What else had changed? The terms "right" and "economic" appeared 32 and 7 times respectively in the 1980 Guidelines. They appeared 61 times and 48 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 over the last two 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 California Consumer Privacy Act (CPPA) in the US and General Data Protection Regulation (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.

Colleague economists may disagree with my interpretation of the trends of the privacy debate. They may lament - much like Posner four decades ago - regulators' archaic reliance on value-based arguments and their blindness to the soundness and objectivity of economic arguments. Some may even consider what I

[15] See https://www.oecd.org/sti/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflows ofpersonaldata.htm
[16] See https://www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf

detect as an emergent unintended consequence a very much intended, and very well needed, progression in the policy discussion around privacy. Yet, if value-grounded arguments had remained so powerful and persuasive among policymakers, US regulators would have eventually 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. 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 General Data Protection Regulation (Atikcan and Chalmers 2019; Christou and Rashid 2021).

Considering the vast network of organizations lobbying against privacy regulation, and 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 plausible: 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. An emerging (and worrisome) policy mindset would be that, if there is no quantifiable economic harm, then there is no privacy concern worth worrying about.

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 problematic, instead, if we do not exercise similar restraint in delimiting - carefully and publicly - also 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" (emphasis added). After focusing his analysis on that one aspect, Posner ended the piece on much broader and broadly assured notes: "[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." 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 *data* protection regulation may have on those benefits are then analyzed; performative, perfunctory references to privacy's other dimensions are interjected, somewhere; and yet broad, encompassing warnings to regulators, with pleads to carefully consider the unintended consequences of their interventions, are the offered conclusions.

As economists we are certainly permitted to articulate the implications of our research.[17] What we should be wary of is the risk of an intellectual sleight of hand: studying a part (the effects on a subsets of directly

[17] It is 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 am more sanguine about the latter and about

measurable, hand-picked metrics) but making conclusions for a whole (broad warnings to regulators) which our analyses have barely grazed.

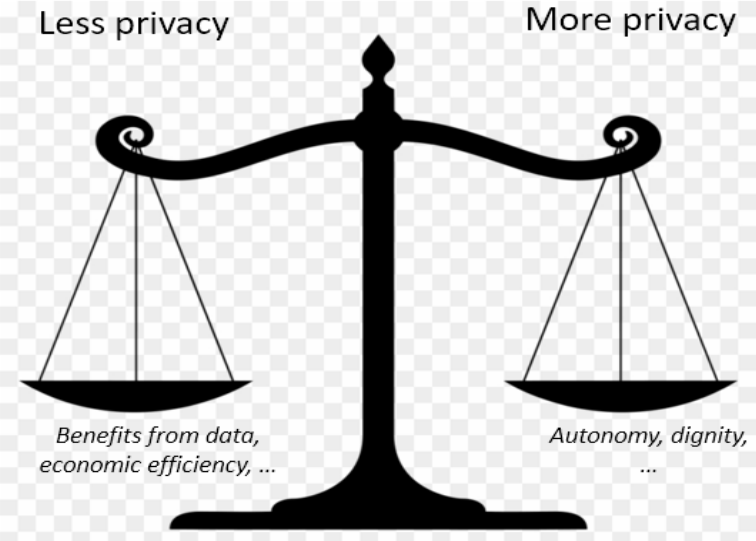

Figure 1

Section 3: Turning the tables and 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 captures 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. It then proposes alternative ways of framing economic research around privacy, and suggests research questions that are complementary to the current focus on the costs of privacy regulation. Finally, it proposes a radical re-framing of the economic debate around privacy that turns the table around the expectations of evidence in the public and policy debate.

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. The scale measures and compares benefits to humanity of those two outcomes. Inherent to that framing is the assumption that interventions such as regulation aimed at protecting privacy may increase freedom or autonomy (which are measured on the right-side pan) but may threaten the economic benefits from data (which are measured on left-side pan).

---

articulating the policy implications of available research that addresses it: if a policymaker's goal was to help consumers better manage their privacy, then results from behavioral research do offer clear indications about the (in)capacity of self-regulatory notice and consent mechanisms to allow proper privacy management (see, e.g., Acquisti et al 2013).

Viceversa, refraining from regulating privacy may harm freedom and autonomy but may allow more economic benefits to be extracted from data.

Which side carries more weight - that is, demonstrates more benefits? Under the economic framing of the privacy debate, less privacy means more data, and more data means more tangible economic benefits. Those benefits (left pan) appear tangible and measurable: more free content and services, more innovation, more efficiency, and so forth. In comparison, the benefits of more privacy (right pan) appear valuable but more abstract, more vague. Thus, an economic framing of the debate can nudge us to think that the benefits of less privacy may outweigh the costs, and (viceversa) that privacy regulation may curtail the benefits of data in exchange for little clear reward.

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 antithetical to the extraction of 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. The section concludes by proposing alternative ways of framing economic research around privacy, and by suggesting research questions that are complementary to the current focus on the costs of privacy regulation.

3.1 Missing the Forest for the Tree: 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 rather than uncritically positing that value, we ought to separate the snake oil of analytics from its demonstrable gains, and scrutinize the allocation of those benefits. Extant 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 collected and analyzed. However, historically, it played a key role in the process through which the Internet became an architecture of commercial surveillance, and remains nowadays a primary channel through which consumer data is funneled into a black box of 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."*[18]

The notion of behavioral advertising as an economic win-win for multiple stakeholders has been internalized in some of the academic marketing 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) 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, Facebook, and other stakeholders in the ecosystem) play the role of matchmakers. They use the vast amount of consumer and merchants 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 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-wins for all stakeholders in the market.

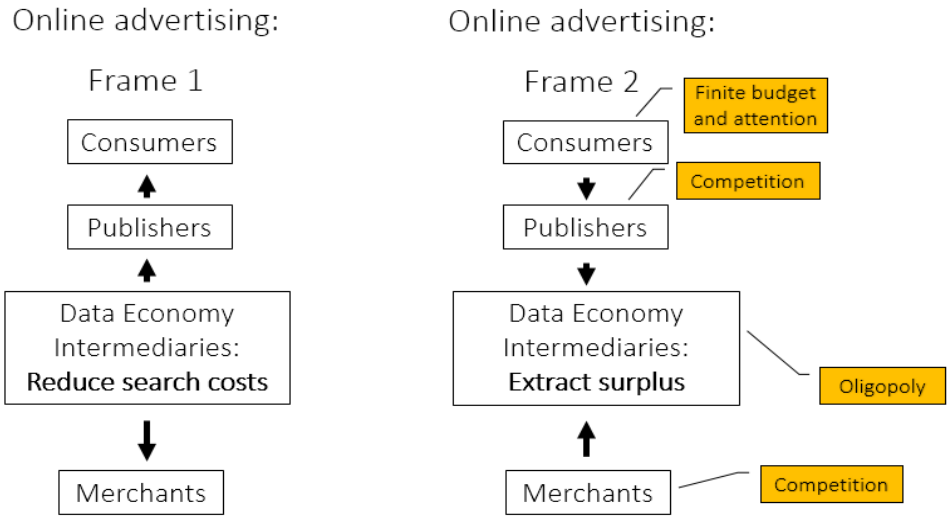

Figure 2

The right side of Figure 2 (Frame 2) presents an alternative economic representation of the online advertising market. The stakeholders are the same. The focus, however, changes from search costs to competition, and from the role of intermediaries in reducing those 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.

[18] See https://www.adexchanger.com/online-advertising/why-is-tracking-good/.

Central to Frame 2 is the observation that consumers have a finite budget 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 has acted as a double-edged sword for publishers in two ways. 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 traditional publishers (at any moment, an online publisher - for instance, nyt.com - may be competing for 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); this is one of the possible reasons why traditional online publishers' revenues have kept decreasing in recent years.[19] 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 more valuable (as targeting is correlated with higher ad conversion rates; Boerman, Kruikemeier, and Zuiderveen Borgesius, 2017) and therefore more profitable for publishers. On the other hand, behavioral targeting has diminished publishers' power in the matching of consumers with advertisers, and thus arguably reducing publishers' revenues. It is important to understand why: before the advent 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 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). This stretches out the supply of advertising spaces - the set of outlets and channels where merchants can find (and buy ad spaces for) interested consumers, and shifts the power to match consumers with advertisers away from publishers and towards third party data intermediaries. With that shift, the ability to extract surplus from advertising transactions also shifts from publishers to intermediaries. This is the second potential reason why traditional publishers' revenues have kept decreasing notwithstanding (or 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 advent of behavioral advertising, a merchant selling golf-related products intent on advertising to golf-interested consumers may have allocated its advertising budget to affinity 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). And while online advertising has also supported a

[19] See https://www.pewresearch.org/journalism/fact-sheet/newspapers/.

proliferation of content distribution channels, the resulting increase in ad spaces through which merchants can reach potential customers does not lessen competition across advertisers, because there is a finite upper boundary to how many ads a consumer can pay attention to and how many products she can buy.

Finally, while the evolution of online tracking and targeting has created more competition on both sides of the market depicted under Frame 2 in Figure 2, it has allowed - in the middle of it - an ecosystem of intermediaries dominated by few large oligopolies. Under this market structure, the oligopoly firms in the middle of the market may be able to extract more surplus from advertising transactions than the aggressively competing stakeholders on either side. Under this framing of the economic impact of online advertising, the intermediaries are those ultimately benefiting the most from the data economy.

Both Frame 1 and Frame 2 of Figure 2 are based on plausible theoretical arguments (Bergemann and Bonatti 2022 highlight how digital platforms can generate both dynamics: surplus creation from matching, and surplus extraction from market power). To some extent, both the search cost reduction story and the oligopoly intermediaries surplus extract story may be in fact occurring. But which side has more empirical validation?

The degree of attention scholarly research has paid to the different stakeholders in Figure 2 is uneven. Oligopoly data intermediaries' record-high profits are evident - although the evidence tends to come more often from industry reports rather than empirical scholarly work.[20] Advertising merchants have received the vast majority of research attention, as a substantial amount of work has examined online advertising effectiveness (Boerman et al 2017). Works in this area have consistently shown that behaviorally targeted ads command higher click-through and higher conversion rates than non-targeted ones. On first analysis, this may support the notion that tracking and targeting technologies help advertising merchants. In reality, their impact on merchants' welfare is almost surely more nuanced than what click-through or conversion rates can capture. As most merchants today can engage in this form of advertising, they may, collectively, wind up at zero-sum prisoner dilemma dynamics. Individually, each advertiser experiences a high conversion rate from behaviorally targeted ads. However, each advertiser has to pay a premium to behaviorally target consumers, merely to avoid competing merchants from poaching consumers away. Hence its market share does not grow: in equilibrium, each advertiser maintains its market share - but spends more to keep it than before. Alternatively, rather than generating prisoner 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 online advertising.

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

[20] 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).

increase revenues for publishers relative to non-behaviorally targeted ones. The increase varies across studies: from over 50% in a study by Google (Ravichandran and Korula 2019), to half of that, in an independent study (Laub, Miller, and Skiera 2022), to a fraction of that in a study using an empirical approach similar to Laub et al (2022), but data from a single large, and arguably sophisticated, media company (Marotta, Abhishek, and Acquisti 2019). As in the case of empirical studies of privacy regulation, these studies offer, individually, 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 tree. 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, but are mute on the *overall* revenue-decreasing effect of the infinite inventory problem. Behavioral advertising giveth, and behavioral advertising taketh away - and yet, to our knowledge, no study has quantified and compared the two contrasting effects. 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 (Apple ATT: users of iOS devices; GDPR: EU residents that 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). Therefore, they do not impair the tracking and targeting of many other categories of users, who remain trackable and targetable. Hence, those studies arguably capture the 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. 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 correlational evidence, coming from publishers' balance sheets, of continuous declines in revenues associated with the advent of behavioral advertising.

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, however, is also limited, and ultimately inconclusive. Search costs are but one factor in consumer utility. Other factors that will surely 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 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 some of that counterfactual evidence has started emerging. In a recent working paper, we found that purchasing products from ads, rather than from search results, increases the likelihood of purchasing from a lower quality merchant and increases the expected price of the product (Mustri, Adjerid, and Acquisti 2020). 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 costs 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 general, in supporting the provision of content and services is indisputable. Many online services *are* supported via ads. The direct impact of those ads on consumer welfare is unclear (see above), but many consumers are comfortable "paying" for online services with their eyeballs rather than with cash (although a substantial amount of consumers now prefers to block ads altogether).[21] The role of *behaviorally* targeted ads in particular 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-edge effect that, we noted above, behavioral advertising can have on the revenues of content creators. I argue that, 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 advent of behavioral advertising in (roughly) the mid-2000s. At the time, those services and 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?[22] Asked differently: to what degree is the connection between increase in behavioral data collection and increase or improvement in online content causal, rather than merely correlational? To what extent is the degradation of privacy an unavoidable and necessary price to pay for more or better content? Is that degradation a necessary condition for innovation?[23]

Conceptually, these questions amount to simple economic comparisons of the marginal cost of privacy loss and the marginal benefits of data collected. Empirically, answering those questions is anything but simple. We face an array of disparate pieces of anecdotal evidence, but limited rigorous work. 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 uses 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'

[21] See https://www.insiderintelligence.com/content/ad-blocking-growth-is-slowing-down-but-not-going-away.
[22] For instance: the number of average ads *per video* has kept increasing on YouTube over time; to what degree has that increase led to more or better YouTube videos or services?
[23] For instance: over the past two decades, Facebook has gained access to more consumer data than most other companies in the world. To what degree has the increase in the amount of identified user data over time led to societally beneficial innovations?

incentives to create new apps (Janssen, Kesler, Kummer, and Waldfogel 2022), but Apple's introduction of ATT does not appear to have negatively affected the supply of new apps for iOS users (Cheyre et al 2022); furthermore, the GDPR does not appear to have negatively affected quantity and quality of EU news and media websites' content (Lefrere et al 2022).

The issue considered in this section was 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 suggests that these are open questions. We know that there are benefits from data, but we do *not* have robust evidence for how the data economy benefits its various stakeholders. In other words, we do not have rigorous economic corroboration for the left side (Frame 1) of Figure 2, nor support for the economic framing that pitches less privacy as a necessary condition for shared benefits from data. And this null result is, in and of itself, a remarkable finding.

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 orthogonal 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 have materialized in the US at both the federal and state levels. And yet, those regulatory efforts did not 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 (almost) trillion-dollar data economy, the growth of data-driven new products and services, and record profits for several intermediaries. Is there a disproportion between economists' fears about, and the actual impact of, privacy protection? Are privacy and analytics (and the extraction of value from data) inherently antithetical, or could both of them be simultaneously achieved through a combination of technology and targeted policy intervention?

As we noted in Section 2, several careful studies have 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, Acquisti, Telang, Padman, and Adler-Milstein 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: the direction of the impact will vary based on how interventions are designed, implemented, and enforced (BERR 2008).

Second (and, as usual, with exceptions: see, for instance, Janssen 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 small and even null effects (several possible explanations exist: see Lefrere et al 2022).

Third (and, again, with exceptions: consider for instance 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 downstream welfare effects on consumers (for instance, a reduction in consumer welfare due to less precisely targeted ads, or a reduction in their usage or satisfaction with available apps).

Fourth, 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 therefore will 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 include 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 downstream long-term effects of increased consumer protection and competition and innovation in privacy preserving analytics between firms.

Sixth (and related to the last point), 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 or homomorphic encryption— that make it possible to analyze and extract value from data while protecting privacy. Granted, there is no free lunch: reducing the granularity of data can be costly, as it also reduces its value. But research suggests that those costs may be carefully minimized (Abowd and Schmutte 2019). In recent work, we considered how the application of differentially private mechanisms to Census data affects educational funding calculations (Steed, Liu, Wu, and Acquisti 2022). 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.

In short, looking at the privacy regulation economic literature in context casts doubts over the notion of privacy protection as being systematically welfare-destroying and inherently antithetical to data utility. Not only is the available evidence on the impact of privacy regulation both incomplete and mixed, but privacy technologies suggest that privacy protection may not need to conflict with data analytics.

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 costs of privacy invasions, and has more often focused on capturing the costs of regulatory intervention. Such focus bore two consequences. First, economic research has produced an inadequate amount of evidence on the harms of less privacy and the benefits of more privacy. This creates a knowledge gap that humpers evidence-based policy-making. Second, 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 versus no less important but less salient, less direct, and less tangible factors. The framing therefore emphasizes the importance of one side over the other.[24]

The scale 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 fundamental *functions* privacy plays within a society, and thus the real - albeit sometimes intangible and hard to quantify - repercussions that privacy dynamics have on individuals' and societies' welfare. The scale misses the economic dark matter - the vast evidence of privacy harm we discussed in Section 2 and exemplified through four scenarios. The incipits of those scenarios were informational: they all involved data collection and usage. Their immediate consequences seemed, at first glance, intangible and abstract: loss of bodily or decisional autonomy. But in fact, we showed that their downstream repercussions on individuals' and societal welfare were real. In other words, the scale is economically myopic: it focuses on the tip of the iceberg (the costs of protection) while sidestepping the mass of the iceberg itself - the massive implications and repercussions (including economic repercussions) of loss of control over personal data.

Whether it is prudent or advisable to measure the economic dark matter is a valid question. The wisdom of considering certain values untradeable precisely lies in the knowledge that those values are essential to the functioning of a society even though they may not be (on first analysis) economically efficient or measurable.[25] Policymakers (and, more broadly, the public debate around privacy) are therefore stuck in a seemingly unresolvable dilemma. On the one hand, policymakers 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). And if we embraced that threshold, we would have to conclude that our economic research is currently failing policymakers, because it is not measuring net harms. On the other hand, the

[24] 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). A similar challenge is well known in the computer security literature: the money spent to secure information systems is tangible; the impact of those investments (the reduction in potential future damages of security failures) is less tangible, harder to quantify, and uncertain.

[25] Discussing when evidence of privacy harm should be required in privacy lawsuits, Citron and Solove (2022) offer that harm should not be required "because it is irrelevant to the purpose of the lawsuit."

challenge of measuring net harm may in fact prove ultimately insurmountable. We considered, in Section 2.2 some of the hurdles: the aggregation problem, the unbounded downstream effects, the intermingle of positive and negative externalities of data.[26]

But if so, what should guide policymakers? Perhaps, the evidence we presented across this manuscript. The evidence that consumers care for privacy, act to protect it, and want more of it. The evidence that economic and behavioral hurdles make it infeasible for individuals to adequately manage their privacy online. The evidence that the costs of regulation may be overblown in the current debate, and that there are, today, tools to allow data analytics and privacy protection. The evidence that proof of a fair allocation of benefits from data is scant. And the evidence that economic research has bypassed a massive amount of privacy costs.

3.4 Changing the Frame of the Privacy Economic Debate

By attacking privacy regulation as an example of perverse government intervention, dismissing key dimensions of privacy as extraneous to the economic debate around privacy, and embracing the value of unfettered consumer data collection, Posner defined a way of framing the economic debate around privacy that still influences our field today. So far, in this section, I have used an economic perspective to highlight systemic problems with that framing. 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.

If my critique has merit, we may then as well alter the frame of the economic privacy debate. Rather than uncritically accepting that "Privacy protection is often costly and at worst inefficient; unless one can demonstrate quantifiable privacy harms, what need is there for govt regulation?," we could ask instead: "What is the evidence that current product and services cannot be provided in more privacy-preserving manners?" Rather than focusing on "What is the evidence of net privacy harm to consumers?," we could instead ask ourselves: "What is the evidence that firms cannot provide current products and services in a more privacy preserving manner?" This is, in essence, a call for turning the tables in the economic privacy debate.

To further that debate, we need to better understand the nuanced and complex interplays between privacy and economic value. We need to foster the type of research (the "exceptions" I have cited throughout the manuscript) that tackles less studied research questions. We need to better understand the harm of privacy loss: Can we (and should we) calculate the economic dark matter? How do we tackle the "aggregation" problem of privacy harms? How do we help consumers and policymakers process the current asymmetry

[26] Data uses will give rise to both positive and negative externalities for individuals and for society at large. For instance, trackers on websites (Scenario 1) may allow targeting of ads that a proportion of consumers may value. Since privacy is not about restricting information flows but managing them, privacy tools (including privacy preserving analytics and well-designed policy interventions) aim at preserving positive externalities (in terms of both individual and aggregate welfare) and curtailing negative ones. In some contexts (such as price discrimination), the definition of positive and negative will be relative to the viewpoints of the data subject and the data holder. As the economic interests of different stakeholders may conflict over the desired balance of collection and protection, weighting those different interests is an ineludible task in policymaking (in privacy, there exists no neutral default setting). Whether a policymaker chooses or not to intervene in disputes over personal data, the policymaker *is* actively making redistributive decisions. See Section 2 and the comment on Noam's analysis.

being tangible benefits of data and intangible harms of privacy? 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 their beneficial effects? What are the economic effects of the deployment of privacy enhancing technologies? Who bears their costs? What are their allocative effects? When we deploy differential privacy, are we decreasing quality of service for consumers? Are we actually burdening society as a whole by making it harder or even impossible to develop new technologies? Or are we decreasing the profit margins of data oligopolies? And we need to better understand the allocation of value from data: How is the value of data allocated? Ultimately, who benefits from the data economy?

Section 4: Conclusions

The debate we considered in this manuscript started over forty years ago. As Posner (1978) decided to "sidestep the definitional problem by simply noting that one aspect of privacy is 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." 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," 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."

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 argued, the economics of privacy ultimately risks crowding critical dimensions of privacy not merely out of its own field of research, but out of 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 flows today, the intellectual continent of privacy has become possibly even vaster than Hishleifer 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 theorizing from other social sciences, and by admitting that a drive for privacy is not antithetical to the extraction of benefits from data - since we have both technologies and strategies to 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 a third way, which consists in focusing on a different set of research questions 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. Time will tell which path it will take.

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.

Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L.F. and Agarwal, Y., 2015, April. Your location has been shared 5,398 times! A field study on mobile app privacy nudging. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems* (pp. 787-796).

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.

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.

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.*

BERR (Department for Business, Enterprise, and Regulatory Reform), 2008. Regulation and innovation: Evidence and policy implications. BERR Economics Paper no. 4. BERR, UK.

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., Kruikemeier, S. and Zuiderveen Borgesius, F.J., 2017. Online behavioral advertising: A literature review and research agenda. *Journal of advertising*, *46*(3), pp.363-376.

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.

Buckman, J.R., Adjerid, I. and Tucker, C., 2022. Privacy Regulation and Barriers to Public Health. *Management Science*.

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.

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*.

Farrell, J., 2012. Can privacy be just another good. *J. on Telecomm. & High Tech. L.*, *10*, p.251.

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., 2019. Regulating privacy online: The early impact of the GDPR on European web traffic & e-commerce outcomes. *Available at SSRN*, *3421731*.

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*.

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*.

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*.

Jones, C.I. and Tonetti, C., 2020. Nonrivalry and the Economics of Data. *American Economic Review*, *110*(9), pp.2819-58.

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.

Miller, A.R. and Tucker, C.E., 2011. Can health care information technology save babies?. *Journal of Political Economy*, *119*(2), pp.289-324.

Mustri, E.A.S., Adjerid, I. and Acquisti, A., 2020. Behavioral advertising and consumer welfare: An empirical investigation. *Working Paper presented at WISE*.

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.

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.

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.

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.

Seeman, J. and Susser, D., 2022. Between Privacy and Utility: On Differential Privacy in Theory and Practice. *Available at SSRN 4283836*.

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.

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.

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.

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

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

Varian, H.R., 1996. Economic aspects of personal privacy, privacy and self-regulation in the information age. *National Telecommunications and Information Administration Report*.

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

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