Bank of Japan Working Paper Series

# The Economics of Privacy: A Primer Especially for Policymakers

Yosuke Uno*
yousuke.uno@boj.or.jp

Akira Sonoda**

Masaki Bessho***
masaki.bessho@boj.or.jp

No.21-E-11
August 2021

Bank of Japan
2-1-1 Nihonbashi-Hongokucho, Chuo-ku, Tokyo 103-0021, Japan

* Payment and Settlement Systems Department (currently at Financial System and Bank Examination Department), Bank of Japan

** Bank of Japan

*** Payment and Settlement Systems Department, Bank of Japan

# The Economics of Privacy*

## A Primer Especially for Policymakers

Yosuke Uno† Akira Sonoda‡ Masaki Bessho§

August 6, 2021

**Abstract**

This paper presents a survey of a field called the *economics of privacy*. Reflecting growing concerns worldwide about the handling of personal data on the Internet, the economics of privacy is developing rapidly, coinciding with recent efforts by privacy regulators to tighten regulations. The literature argues that it is difficult for market mechanisms to resolve problems such as how to determine the socially optimal level of privacy protection and how to avoid excessive privacy loss driven by *negative data externalities*. These insights should be useful for policymakers facing the question of how to deal with personal data issues and to ensure that people's privacy is protected. (*JEL* D62, D82, D83, K20, M31, M37)

# 1 Introduction

The protection of privacy on the Internet has become a major global concern in recent years. One of the triggers was the *Cambridge Analytica* scandal. *Cambridge Analytica* obtained the personal data of 50 million Facebook users and used it for personalized political advertisements and/or the planting of "fake news" in the

*We thank Yasutora Watanabe, Kazushige Kamiyama, Yutaka Soejima, Akio Okuno, Takeshi Yamada, Junichiro Hatogai, Masashi Hojo, Masashi Une, Kazutoshi Kan, Noriyuki Shiraki, and Hiroyuki Takano for useful comments. All remaining errors are ours. This paper does not necessarily reflect the views of the Bank of Japan.

†Bank of Japan (yousuke.uno@boj.or.jp)

‡Bank of Japan

§Bank of Japan (masaki.bessho@boj.or.jp)

2016 US presidential election and the UK referendum on leaving the European Union.[1,2]

This scandal has made people aware of the vast amount of personal data that online platforms hold and the huge impact the improper use of such data can have. At the same time, the business model of online platforms that collect and monetize vast amounts of personal data has also come under scrutiny.[3] For example, social network services users post photos with text and tags to communicate with their friends, but behind the scenes, providers use those photos as *labeled image data* to train artificial intelligence (AI) algorithms. This system is sometimes referred to as *technofeudalism* because it can be seen as a structure where online platforms provide useful and enjoyable information services, while taking all the upside profit of the data consumers create in exchange (Posner and Weyl 2018).[4] Lanier (2013) worries about the social and economic consequences of online platforms' business model since they do not give their users proper incentives to supply their personal data.

Against this backdrop, regulators in major jurisdictions have introduced a series of privacy regulations in recent years. For instance, the European Union (EU) in May 2018 introduced the General Data Protection Regulation (GDPR), which requires the informed consent of data subjects to the processing of data. Europe has a long history of privacy regulation, and the GDPR was created as a successor to the EU Data Protection Directive, which was passed in 1995. In California, the California Consumer Privacy Act of 2018 (CCPA), which strictly regulates the handling of consumers' personal information, went into effect in January 2020.

In line with these regulatory developments, researchers have been actively discussing how to understand the behavior of online platform providers, how to protect people's privacy, and how privacy regulations such as the GDPR and CCPA can work. Knowledge of areas such as machine learning and computer science is essential to understand the business of online platforms. Knowledge of law is also

[1]The Guardian, "How Cambridge Analytica Turned Facebook 'Likes' into a Lucrative Political Tool," 17 March 2018; New York Times, "How Trump Consultants Exploited the Facebook Data of Millions," 17 March 2018; BBC, "Cambridge Analytica Planted Fake News," 20 March 2018.

[2]With regard to the 2016 US presidential election, there are various views on whether personalized political advertisements or "fake news" on social media influenced the outcome of the election. For example, Allcott and Gentzkow (2017) suggest that "fake news" on social media did *not* affect the outcome of the election.

[3]Online platforms track people's behavior on the Internet. The six companies that are most actively engaged in such activities are Google, Facebook, Twitter, Amazon, AdNexus, and Oracle (Englehardt and Narayanan 2016).

[4]In order to address the *feudal* state, Posner and Weyl (2018) and Arrieta-Ibarra et al. (2018) propose the idea of *data as labor*, which considers personal data created on the Internet as a product of labor.

required to understand the basic concepts of regulations such as the GDPR and CCPA. In addition, it is useful to consider relevant issues from an economics perspective in order to understand various trade-offs. The *economics of privacy* use knowledge and tools from these different domains to discuss what forms privacy protection should take.[5]

This paper reviews some of the insights of the economics of privacy and should be useful for policymakers facing the question of how to deal with personal data issues and to ensure that people's privacy is protected. For example, payment processing involves the transfer of certain personal data (sender, recipient, value, date, and time). In recent years, the digitization of payments and the diversification of payment interfaces have led to new forms of personal data transfers within the ecosystems in and around payments. Online platform providers are entering digital payment services in search of new personal data (FSB 2019). Protecting people's privacy likely would be one of the most important issues in the design of payment systems. In fact, in the European Central Bank's public consultation on a digital euro in 2020, privacy would be the most important feature of a digital euro, mentioned by 43% of respondents to a survey (ECB 2021). This means that the protection of privacy forms the premise of any discussion on how firms could use the data created in a digital payment system.

The remainder of this paper is organized as follows. Section 2 lays out the basics of privacy. Section 3 discusses the cost of privacy protection paid by firms. While privacy protection provides benefits to consumers, it also imposes costs through restrictions on the use of personal data. This cost is paid by firms that monetize the use of personal data. In the economics of privacy, however, some studies interestingly document that protecting consumer privacy also brings benefits to firms. In Section 4, we focus on *negative data externalities*. Negative externalities arise in situations where information that a person wishes to keep secret can be inferred by prediction. The implications of this for privacy protection are serious, and it is the most important issue in the economics of privacy. Finally, Section 5 summarizes this paper.

[5] According to Acquisti et al. (2016), the current surge in interest in the economics of privacy is the third wave. The first wave was driven by the *Chicago School* in the 1970s and 1980s, which included scholars such as Posner (1978, 1981) and Stigler (1980) and argued that privacy protection creates inefficiencies because it hides useful information. The second wave, in the mid-1990s, discussed the role of cryptography and the implications of secondary uses of personal data. For example, Varian (1996) argued that the costs incurred by consumers could be significant if *too little* personal data is provided to third parties.

# 2 Foundations of Privacy

This section reviews basic issues related to privacy. Specifically, after looking at how economists often think about privacy (Section 2.1), we review the concept of *differential privacy*, which is a formal quantitative framework for guaranteeing privacy protection (Section 2.2). Using differential privacy as a tool allows us to conduct various policy discussions. However, as we will see in Section 2.3, observing the differential privacy parameter involves numerous hurdles, and, as will be discussed in Section 2.4, designing rational differentially private systems via market mechanisms may be challenging.

## 2.1 What is Privacy?

Privacy is difficult to define, which means that privacy means different things to different people (Posner 1978, 1981; Acquisti et al. 2016). Some may consider information on their employment status to be private, while others may regard information related to their health as private. Moreover, even for the same person, privacy is likely to be context specific. For example, a person may regard information that they went to a restaurant for dinner as privacy information depending on who they went for dinner with.

The economics of privacy often considers privacy to have *instrumental* rather than *intrinsic* value to people's utility (Posner 1978, 1981; Acquisti et al. 2016). For example, the literature regards the value of privacy as instrumentally created when personal data that contains privacy information is used for marketing purposes. As we will see in Section 2.3.1, a recent study by Lin (2021) attempts to clarify the instrumental and intrinsic aspects of privacy through a well-designed experiment.

As will be discussed in more detail in Section 4, a major concern of the economics of privacy is negative data externalities, which implicitly assumes that what people consider as privacy information differs. Note that, in some cases, such as in macroeconomic discussions that assume a representative consumer, differences in what people consider as privacy information are not assumed (e.g., Jones and Tonetti 2020).

## 2.2 Differential Privacy: A Tool

In the economics of privacy, when describing the degree of privacy protection, a tool called *differential privacy* is often used (Ghosh and Roth 2011; Pai and Roth 2013; Hsu et al. 2014; Abowd and Schmutte 2019). This subsection provides an

overview of differential privacy (Section 2.2.1) and then incorporates the concept of differential privacy into a utility function for privacy (Section 2.2.2).

### 2.2.1 Differential Privacy

Differential privacy is a well-known concept in the field of computer science (Dwork et al. 2006; Dwork 2006). Let $D$ be a dataset containing personal data, and a query action to retrieve information from the dataset $D$ is denoted by $Q$. In that case, even if the query does not identify individuals, the result of the query $Q(D)$ may reveal personal information contained in the dataset $D$. For illustration, imagine that we execute a query to extract "gender of the examinee" and "pass/fail" results from a dataset that records the exam results of a class of 20 males and 20 females. If the query is run and all successful candidates are female, personal information of males in the class who wanted to keep their exam results secret is inadvertently divulged and their privacy has been breached, since the query result revealed that all males failed.

Now, instead of query results being output as they are, we consider them being output after some processing to protect privacy. Let $\mathcal{M}$ and $\mathcal{M}(D)$ be privacy-protection processing and the query result after the privacy-protection processing, respectively. How can we evaluate the reliability of the privacy-protection processing denoted by $\mathcal{M}$? The concept of differential privacy addresses this question by considering two datasets, $D_1 \in D$ and $D_2 \in D$, which differ only in whether they contain the data of one person, and assesses the reliability of $\mathcal{M}$ based on whether or not some information about the person can be obtained from the *difference* between the query results $\mathcal{M}(D_1)$ and $\mathcal{M}(D_2)$.

Let us use a real-world example to explain. Let $D_1$ be a dataset that contains the salary of the staff members of a certain department at period $t$. At period $t+1$, a new staff member is hired. Let $D_2$ be a dataset contains the salary of the staff members of the department at period $t+1$. We consider a query for the average salary, denoted by $Q$. Let $\mathcal{M}$ be a procedure that outputs $Q(D_1)$ and $Q(D_2)$ after adding random noise to the datasets $D_1$ and $D_2$.

Running the query $Q$ 100 times for $D_1$ and $D_2$ provides us with 100 observations of the noisy average salaries $\mathcal{M}(D_1)$ and $\mathcal{M}(D_2)$. We can then examine the shape of the two distributions of observations of the average salaries $\mathcal{M}(D_1)$ and $\mathcal{M}(D_2)$. If the two distributions match perfectly, the new employee's privacy can be regarded as being completely protected, since the distributions reveal no information about her/his salary. However, such complete privacy protection also means that we can obtain no additionally useful information from the dataset $D_2$.

We now consider the trade-off between protecting the privacy of the newly hired employee and ensuring the usefulness of the dataset $D_2$. The breakthrough idea of differential privacy is that it expresses this trade-off with only one parameter, $\varepsilon$. Formally, a privacy protection scheme $\mathcal{M}$ satisfies differential privacy if for all pairs of *neighboring* datasets $(D_1, D_2)$ differing only in one person's data, and for all $R \subseteq Range(\mathcal{M})$,

$$\frac{\Pr(\mathcal{M}(D_1) \in R)}{\Pr(\mathcal{M}(D_2) \in R)} \leq \exp(\varepsilon)$$

for $\varepsilon > 0$.

Intuitively, $\varepsilon$ expresses the size of the discrepancy between the two probability distributions $\Pr(\mathcal{M}(D_1))$ and $\Pr(\mathcal{M}(D_2))$. If $\varepsilon$ is large, the two probability distributions are far apart, and we can obtain some information about the salary of the newly hired employee. A large $\varepsilon$ is desirable when the priority is to ensure that the dataset $D_2$ is useful. On the other hand, the smaller $\varepsilon$, the more indistinguishable are the two probability distributions. Therefore, a small $\varepsilon$ (of close to zero) is desirable if the priority is privacy protection.[6] Note that $\varepsilon$ is referred to as the privacy loss or privacy budget.

Unfortunately, it is not possible to gain intuition about how much privacy is protected if, for example, $\varepsilon = 0.01$. Importantly, the idea of differential privacy provides a quantitative measure of privacy protection instead of a binary assessment of the question as to whether privacy is protected or not.[7]

### 2.2.2 Differential Privacy in a Utility Function

Ghosh and Roth (2011) use the parameter of differential privacy, $\varepsilon$, to specify the following utility function with respect to privacy:

$$u_i = p_i - v_i\varepsilon \tag{1}$$

[6] In this discussion, it is implicitly assumed that the dataset administrator can access the personal data before the noise is added. However, some people may not want to disclose their personal data even to the dataset administrator. To address this situation, a technique called *local differential privacy (LDP)* has been proposed, which guarantees stricter privacy protection by not disclosing personal data even to the dataset administrator (Kasiviswanathan et al. 2011; Duchi et al. 2013). See Appendix A for more details.

[7] Without the use of quantitative metrics such as differential privacy, it is risky to assume intuitively that if certain personal information, such as addresses and phone numbers, were removed from a dataset, it would not be possible to identify individuals. In computer science, a few cautionary tales that illustrate how intuition-based attempts at anonymization have failed are well known (Narayanan and Shmatikov 2008; Heffetz and Ligett 2014).

where $u_i$ denotes the utility of consumer $i$, $p_i$ denotes the compensation for a loss of privacy for consumer $i$, $v_i$ represents the disutility from the loss of privacy of consumer $i$, and $\varepsilon$ is the privacy budget. Note that $p_i$ is not necessarily monetary and includes the convenience provided by applications or the enjoyment derived from the use of online services.

## 2.3 The Difficulty of Observing $\varepsilon$

Observing $\varepsilon$ in real-world data is difficult in two regards. The first is that it is difficult to obtain sufficient information on $u_i$, $p_i$, and $v_i$ in (1) to identify $\varepsilon$. For example, online services that provide navigation on a map are usually provided in exchange for the user's current location data. Assuming that a dataset to observe whether the service is being used is available, since the provision of the service and the provision of the users' personal data occur simultaneously, it is not possible to identify whether users use the service since they are not very concerned about the provision of location data, i.e., $v_i\varepsilon$ in (1) is small, or because they highly value the convenience of the service, i.e., $p_i$ in (1) is large. Despite these difficulties, as we will see in Section 2.3.1, a number of studies have attempted to value the utility of privacy.

The second respect in which it is difficult to observe $\varepsilon$ in real-world data is that there is a discrepancy between the desired degree of privacy protection that people express in surveys and their actual behavior. In the field of computer science, this is known as the *privacy paradox* (Acquisti 2004; Barnes 2006). We review this paradox in Section 2.3.2.

### 2.3.1 Valuing Privacy

Huberman et al. (2005) conduct a reverse second-price auction, asking participants how much they would be willing to accept in exchange for personal data such as their age and weight. The bidding prices allow the auctioneer, i.e., the authors, to observe $p_i$ in (1). They find a large variation in $p_i$. Goldfarb and Tucker (2012a) attempt to measure $v_i\varepsilon$ in (1). Specifically, they find that over the period 2001 to 2008, people's concerns about privacy, $v_i\varepsilon$, increased year over year, and that older people tended to be more reluctant to disclose information than the young. Moreover, the gap between the two groups widened over the years.

Kummer and Schulte (2019) attempt to observe revealed preferences for privacy from data on about 300,000 smartphone apps observed on the Google Play Store between 2012 and 2014. The first step of their analysis consists of identifying

if the apps contain privacy-sensitive permission information using a feature of the Google Play Store that allows app developers to choose among standardized blocks of information, so-called permissions, where some enable access to a user's location, communication, browsing behavior, etc. The authors then examine whether the presence or absence of privacy permissions affects supply and demand in the *privacy market*. The results of the estimation show that requiring access to privacy-sensitive information reduces the number of installs by 25% on the demand side and significantly decreases the price on the supply side.

Lin (2021) attempts to measure the utility of privacy by conducting well-designed experiments. Specifically, based on Becker's (1980) utility model, the utility of privacy is decomposed into an *intrinsic* part and an *instrumental* part, and Lin (2021) finds that the subjective evaluation of the intrinsic part of privacy varies widely among people. At the same time, she documents that some people have extremely high subjective valuations, i.e., the distribution is skewed to the right.

### 2.3.2 The Privacy Paradox

The apparent dichotomy between privacy concerns and actual privacy behaviors has caught the attention of researchers. This is known as the *privacy paradox*, a phenomenon that while people claim to be very concerned about their privacy, they nevertheless undertake very little to protect their personal data that contains privacy information.

Using data from a social experiment conducted at the Massachusetts Institute of Technology in 2014, Athey et al. (2017) report that a digital privacy paradox was observed.[8] In the experiment, subjects were randomly divided into a control and a treatment group, where those in the treatment group were given a coupon for a free pizza with their closest friends. Both groups were then instructed to give the e-mail addresses of their friends to the experimenter. The result of the experiment showed that the probability that those in the treatment group gave invalid e-mail addresses was 54% lower than in the control group, despite the very small incentive of a free pizza. This result was stable even after taking into account differences in *ex ante* stated preferences about privacy.

Using the concepts of the willingness to pay (WTP), that is, the amount of money a person would be willing to pay to acquire a good he/she did not own, and the willingness to accept (WTA), that is, the amount of money a person would require to be willing to accept to part with a good, Acquisti et al. (2013) conducted an

[8]For details of the experiment, see Catalini and Tucker (2016, 2017).

experiment in a shopping mall in Pittsburgh and found that people's WTP to protect the privacy of their data and their WTA to give up privacy protection differ substantially. Such a gap between the WTP and WTA is a well-known phenomenon in the field of behavioral economics called the *endowment effect*, a bias in which people place a higher value on goods they already own. However, Acquisti et al. (2013) find that for privacy, the ratio of the WTA to the WTP is 5.47, which is much larger than the 2.92 for ordinary private goods.[9]

Recently, some empirical studies have attempted to resolve the privacy paradox. For example, Chen et al. (2021) conducted a survey on privacy concerns among users of Alipay, an online payment platform, and analyzed the relationship between privacy concerns and personal data provision behavior by matching the results of the survey with users' administrative data provided to Alipay. The results show that even after controlling for user characteristics, there is no statistically significant relationship between privacy concerns and personal data provisioning behavior; rather, users with stronger privacy concerns are more likely to actively use digital services. A possible explanation of this seemingly paradoxical result is that learning through the use of digital services (a larger $p_i$ in (1)) leads to greater privacy concerns (a larger $v_i$ in (1)). If $p_i$ and $v_i\varepsilon$ are correlated, this could resolve the privacy paradox that the behavior of providing personal data is not consistent with $v_i\varepsilon$. This interpretation suggests that people's concern for protecting their privacy is not innate, but rather a developed preference that is gradually formed through the use of digital services.

### 2.4 Discussion of Policies with regard to $\varepsilon$

The concept of differential privacy is useful for the discussion of policies on the protection of privacy. Abowd and Schmutte (2019) argue that the socially optimal level of $\varepsilon$ is determined in a trade-off between privacy loss and the statistical accuracy available to national statistical agencies such as the Census Bureau in the United States. Ghosh and Roth (2011) and Hsu et al. (2014) discuss the trade-offs faced by policymakers in more general settings as follows.

Imagine that a policymaker plans to construct a database that contains personal information. The policymaker can control the degree of privacy protection for participants in the database through choosing the level of the differential privacy parameter $\varepsilon$. If the policymaker sets $\varepsilon$ to a value close to zero, participants' privacy will be strictly protected, and no one will hesitate to participate in that

[9]For more on the gap between WTA and WTP for privacy, see the survey by Hui and Png (2006).

database on the basis of privacy concerns. However, since the data in the database contains a large amount of random noise to protect privacy, the value of that data will be low. On the other hand, if the policymaker chooses higher values of $\varepsilon$ (reduces the noise) to increase the value of the data, people with high privacy concerns (a relatively large $v_i$ in (1)) are more likely to leave the database. In such a case, there will be a relatively large number of people with low privacy concerns in the database. As is well known, the data in such database contain *sample selection bias* (Heckman 1979). Thus, in determining $\varepsilon$, policymakers will have to compromise on at least one of the following three aspects: (i) the number of participants in the database, (ii) the value of the data in terms of noise, and (iii) the value of the data in terms of bias.

Under such trade-offs, can policymakers design a market mechanism to rationally determine the optimal level of $\varepsilon$? Ghosh and Roth (2011) point out that there is *no* individually rational mechanism that would allow people to truthfully report the degree of privacy protection they require.[10] Specifically, they argue that the *direct revelation mechanism* does not work when the disutility, i.e., $v_i\varepsilon$ in (1), and the price, i.e., $p_i$ in (1), are correlated. For example, a person who considers information that he/she has an infectious disease to be privacy information will hesitate to bid a high price in a reverse auction where information on whether or not one has an infectious disease is traded. This is because a high bidding price itself would reveal the fact that one has an infectious disease. Thus, when it comes to privacy, the direct revelation mechanism makes it difficult to rationally determine the level of $\varepsilon$.[11]

Meanwhile, Ichihashi (2020c) discusses privacy protection regulation based on a framework of a dynamic game between a consumer and an online platform. Assuming a decreasing marginal privacy cost for the consumer, Ichihashi (2020c) shows that if a policymaker introduces strict privacy protection regulation (setting $\varepsilon$ close to zero), in the short run, the welfare of the consumer improves and the level of activity on the platform increases. As the consumer's level of activity on the platform rises, more personal data is generated and the consumer's marginal privacy cost is further reduced. In the long run, the consumer's level of activity on the platform becomes so high that the consumer eventually loses her/his privacy. This suggests that, in the long run, strict privacy regulations may have perverse

[10] Here, differential privacy is treated as an equilibrium concept or a solution concept. That is, differential privacy makes the outcome approximately independent of any additionally included single agent's data. This point with regard to differential privacy as an equilibrium concept was highlighted by McSherry and Talwar (2007).

[11] The impossibility result of Ghosh and Roth (2011) can be circumvented by introducing certain different settings (see, e.g., Ligett and Roth 2012).

effects and not achieve their intended purpose.[12]

At present, to the best of our knowledge, there is no methodology of rationally determining the optimal level of $\varepsilon$.[13] As noted by Heffetz and Ligett (2014), "the time seems ripe for more economists to join the conversation," and we hope that more knowledge will be accumulated in this field in the future.

# 3 Costs of Privacy Protection

While privacy protection brings benefits to consumers, it also imposes costs through restrictions on the use of personal data. These costs are paid by firms that monetize personal data. In a differentially private system, the lower the level of $\varepsilon$ (the stricter the privacy protection), the higher the cost to firms. Although $\varepsilon$ is not set in actual privacy regulations, the costs paid by firms are observable.

Essentially, firms will pay some of the costs of privacy protection in a variety of ways, as seen in Section 3.1. However, as shown in Sections 3.2 and 3.3, they do not always pay the costs and may also reap benefits, i.e., enjoy *negative* costs.

## 3.1 Costs of Privacy Protection Regulations

Goldfarb and Tucker (2011), using data on online ad campaigns worldwide from 2001 to 2008, show empirically that the EU's Privacy and Electronic Communications Directive (2002/58/EC) decreased the effectiveness of advertising on average by around 65%.

Recently, a growing number of empirical studies have examined the economic impact of the GDPR, which was implemented in May 2018. For instance, using data provided by Adobe, Goldberg et al. (2019) found that the GDPR led to a 9.7% decrease in pageviews of all websites in the EU, and a 4.2% decrease in pageviews and an 8.3% decrease in revenue for e-commerce websites. In addition, Jia et al. (forthcoming), using data on venture investments from January 2014 to April 2019, report that the implementation of the GDPR reduced the number of monthly venture deals by EU ventures compared to their US counterparts by

[12] The key point of Ichihashi's (2020c) argument is that *ex ante* and *ex post* regulation of privacy protection can have different effects. That is, *ex ante* privacy regulations may reduce the welfare of consumers, while *ex post* regulations, such as those that protect consumers' *right to be forgotten*, may increase consumers' welfare, as shown in Section 3.2.

[13] In computer science, $\varepsilon$ is often set in the range of 0.01 to 10, but there is little or no convincing justification for that range (Hsu et al. 2014).

26.1%.[14]

The design of privacy protection regulations can have life-or-death consequences. Using the variation in privacy regulations among US states, Miller and Tucker (2009) show that strong privacy regulations significantly reduce the adoption of Electronic Medical Records (EMR). Further, Miller and Tucker (2011) document that a 10% increase in basic EMR adoption would reduce neonatal mortality rates by 16 deaths per 100,000 live births. In light of these findings, Goldfarb and Tucker (2012b) highlight that privacy policy is interlinked with innovation policy.

### 3.2 Costs of Protecting the *Right to be Forgotten*

As seen in Section 2.3.2, people's behavior regarding privacy protection often seems to be not rational. Given this, it is important that people have the opportunity to revoke their past decisions with regard to personal information. In this regard, it has been pointed out that people have the *right to be forgotten* (Rosen 2012). In this context, the above-mentioned study by Ichihashi (2020c) based on a dynamic game framework between a consumer and an online platform, shows that if consumers' right to be forgotten is protected, consumer welfare increases. Note that the EU's GDPR explicitly protects the right of data subjects to withdraw consent they have given in the past, i.e., the right to be forgotten.

From the standpoint of firms that monetize personal data, personal data has the potential to create greater added value if it is retained and accumulated over a long period of time. This implies that if people's right to be forgotten is protected through some mechanism, the retention period of personal data may be restricted, resulting in a reduction of the benefits from utilizing personal data.

Chiou and Tucker (2017), examining the impact of changes in the retention period of searchers' personal data on search engine search quality, report that no statistically significant impact could be identified.[15] This suggests that past personal data is not very important for search engine retrieval services; in other words, the cost of protecting people's right to be forgotten is small for firms that provide search engine retrieval services.

It should be noted that Chiou and Tucker's (2017) results may be specific to search engines, where new words are searched every day. Nevertheless, given that data

[14] It should be noted that, as suggested by the title of their paper, "The Short-Run Effects of the General Data Protection Regulation on Technology Venture Investment," this result may be temporary.

[15] Search quality here is measured in terms of whether search engine users visit a website or redo their search depending on the search results displayed.

quality has a significant impact on the accuracy of algorithms and that the domain of the data may change over time, it seems somewhat plausible that discarding older data at a certain point in time will not lead to worse results.

## 3.3 *Negative* Costs of Privacy Protection

There are interesting empirical studies showing that privacy protection has *negative* costs, i.e., it provides benefits to firms.

### 3.3.1 Improving Advertising Performance through Simple Privacy Controls

Tucker (2014) provides an empirical analysis of the impact of Facebook's privacy policy change implemented on May 28, 2010, on advertising effectiveness. The data are click-through rates for Facebook users resulting from an advertising campaign on Facebook by an educational non-profit organization in the United States. The data cover 2.5 weeks, including May 28, and the change in privacy policy during this period was unexpected for Facebook users.

Before the May 2010 changes, Facebook's privacy policy was considered very complex, with as many as 170 options that users had to select to manage their privacy settings.[16] The policy change introduced an easy-to-use privacy control interface, reduced the amount of information that was automatically required to be displayed, and also gave users new controls over how their personally identifiable data could be tracked or used by third parties. The change can be interpreted as giving Facebook users much stronger bargaining power over the control of their personal data.

While the changes were predicted to reduce the effectiveness of advertising on Facebook, the results were the opposite. Following the change in the privacy policy, the click-through rate nearly doubled. This implies that an increase in consumers' bargaining power benefits firms in terms of increased advertising performance.

### 3.3.2 A Positive Externality of the GDPR

Aridor et al. (2020) use data from an online travel intermediary to examine the impact of the introduction of the GDPR. They find that the introduction of the

[16] In fact, complex privacy policies are quite common. Ramadorai et al. (2019) develop a *Gunning Fog Index* to measure the complexity of the privacy policies of 4,078 US firms and point out that at least a college degree is required to understand the median level privacy policy.

GDPR resulted in a 12.5% decrease in cookies.[17] This is probably because the introduction of the GDPR has led to an increase in the number of consumers who explicitly refuse to accept cookies.

At the same time, however, there was a surprising 8% increase in trackability of consumers who explicitly agreed to accept cookies after the GDPR was introduced.[18] Aridor et al. (2020) argue that this may be the result of a reduction in noise as consumers who previously used browser-based cookie blocking tools explicitly refused cookies and were therefore missing from the data observed by the firm.

Specifically, browser-based cookie-blocking tools regenerate cookies each time a website is visited, so that the same user appears under multiple cookie identifiers, resulting in noisy data for consumer-specific analyses. On the other hand, if a user explicitly refuses to accept cookies under the GDPR, the user's cookie information will not be sent, thereby reducing the noise in the data. This interesting consequence can be interpreted as a positive externality of the GDPR.

# 4 Negative Data Externalities

Broadly speaking, Sections 2 and 3 both provide discussions of the differential privacy parameter, $\varepsilon$. In contrast, this section addresses privacy protection schemes denoted by $\mathcal{M}$ in Section 2.2.1. Specifically, in this section, we consider a situation in which even if some data are strictly concealed, it is still possible to infer the data from other data. This is a well-known problem referred to as *negative data externality* (Section 4.1). The economics of privacy has reached a consensus on the serious consequences of this negative externality (Section 4.2) and proposes effective privacy protection schemes in the presence of such negative externality (Section 4.3).

Note that the word *negative* here means that data disclosure has a negative impact on consumers' utility. Conversely, data disclosure can also have positive externalities if the impact on consumers' utility is positive. Since this paper focuses on privacy protection, the following sections consider only negative externalities; however, it should be noted that when examining the impact of data disclosure on consumers' utility from a comprehensive perspective, it is necessary to also

[17] A cookie is a text file that is sent from a web server to a consumer's web browser when the consumer visits a website to store information about that consumer.

[18] The authors define trackability in terms of a measure of how many times the same cookie is observed on a given website over a period of time.

discuss positive externalities, as, for example, Ichihashi (2020b) and Fainmesser et al. (2021) do.

### 4.1 What are Negative Data Externalities?

A data externality is an effect that arises from the disclosure of personal data. Negative data externalities arise in situations where data that consumer $i$ considers as privacy information and wants to conceal can be inferred from data that consumer $j$ does not consider as privacy information and discloses. The inference is possible because these data are correlated with each other.

In this situation, consumer $i$ is considered to suffer a loss of privacy due to the externality generated by the data provided by consumer $j$. Note that the externality is not considered in consumer $j$'s decision making. Her/his choice to provide information is guided only by her/his private benefits and costs.

Let us use an example. Suppose that females are more likely to purchase confectionary when they are depressed and a confectionary manufacturer attempts to target advertising by identifying females likely to be depressed based on various general pieces of personal data.[19] For females who consider their health status as privacy information, inferring whether they are depressed or not would lead a loss of privacy. It should be noted that the overall utility, i.e., $u_i$ in (1), could be positive, depending on the compensation for the loss of privacy such as a special confectionery coupon, i.e., $p_i$ in (1).

In the economics of privacy, this type of negative externality is considered to be one of the most important issues.

### 4.2 Consequences of Negative Data Externalities for Privacy Protection

The economics of privacy has shown that when consumers behave rationally in the presence of negative data externalities, serious consequences for privacy protection arise. Specifically, it is known that in the presence of negative data externalities, there will be *excessive* data sharing on online platforms and the price of data will be depressed (Choi et al. 2019; Acemoglu et al. forthcoming; Bergemann et al. 2020; Ichihashi 2020a, 2020b; Fainmesser et al. 2021). This, in fact, describes

[19] Wei et al. (2020) examine what attributes are actually used in Twitter's targeted advertising. According to them, the most used attributes for Twitter targeted ads are language, age, and location. Gender targeting is much less frequently used.

the current situation well, where vast amounts of personal data are provided to some online platforms at extremely low prices or for *free*.[20]

The mechanism is simple. Since consumers know that they are affected by negative data externalities and that they may not be able to keep personal information secret even if they wanted to, it is optimal for them to provide their personal data for a low price. On the other hand, if there were a central social planner that maximizes consumers' utility and online platforms' profits, the provision of personal data would be controlled to the extent that the total amount of consumers' privacy loss is not too large. Therefore, if consumers maximize their utility in a decentralized manner, the amount of personal data provided will be excessive compared to the case where the social planner chooses the total amount of personal data provided. As a result, there will be an excessive amount of privacy loss in the economy as a whole.[21,22]

The consequence of this negative data externality is that consumers will not be able to truthfully express the degree of privacy protection they desire. Consumers will be placed in a situation where it becomes optimal for them not to choose stronger privacy protection, even though they prefer it.

As highlighted by Choi et al. (2019), this situation cannot be resolved by educating consumers. That is, it is not that consumers are unable to anticipate the invasion of privacy, but that it has become optimal for them to provide personal data even when they know that it erodes their privacy.

Also, enhancing competition among online platforms does not necessarily improve this situation. Choi et al. (2019) argue that the same situation would arise even if there were competition, while Acemoglu et al. (forthcoming) find that competition among online platforms in fact may make the situation worse. Acemoglu et al. (forthcoming) argue that in some cases it may even be better for the economy to shut down the personal data market.

[20] Although Ichihashi (2020c) in his analysis focuses on a different mechanism from externalities, he found that, under the assumption of decreasing marginal privacy costs for consumers, in the long run consumers eventually lose their privacy but keep choosing an excessive level of activity on the platform. Also see the discussion in Section 2.4.

[21] The consequences of this negative data externality may explain the *privacy paradox* discussed in Section 2.3.2 (Bergemann et al. 2020). That is, in the presence of negative externalities, the value of each consumer's personal data becomes so low that it is rational to give up personal data even for a very low price, such as a free pizza coupon.

[22] Ichihashi (2020a) points out another consequence of consumers' rational behavior under negative data externalities. He argues that in the presence of negative data externalities, consumer payoffs worsen since the prices of goods they face are increased.

### 4.3 Privacy Protection Schemes under Negative Data Externalities

As discussed in Section 4.2, negative data externalities give rise to *inefficiency* in terms of excessive data sharing. If this type of inefficiency can be reduced, then policy intervention is justified. In this subsection, we review several privacy protection schemes that could help to reduce inefficiency due to negative data externalities.

The first is a *personalized Pigovian tax* that would *internalize* data externalities (Acemoglu et al. forthcoming). Intuitively, the inefficiency of excessive data sharing due to negative data externalities arises from the rational behavior that individual consumers' do not take the costs of such externalities into account. Therefore, the tax should be imposed depending on the correlation structure of users' personal data. That is, consumers whose personal data is more correlated with the personal data of another consumer will be taxed relatively more to weaken the incentive to provide personal data.[23] While this taxation scheme can theoretically restore the first-best allocation achieved in the presence of a social planner, it is impractical. For example, for an online platform with 10 million users, it would be impossible to calculate the optimal tax at any given time based on a huge correlation matrix of 10 million users' personal data.

The second potential privacy protection scheme is an *opt-in regulation without price discrimination* (Choi et al. 2019). Opt-in consent regulation such as the EU's GDPR requires consumers to provide explicit consent in advance of providing data. Choi et al. (2019) observe that the combination of opt-in regulation requiring an opt-in when platforms collect data beyond the socially optimal level and regulation that does not allow price discrimination can lead to a socially optimal outcome because online platforms would be able to collect data only up to the socially optimal level. This second scheme is essentially based on the same kind of idea as the first one in that it imposes a cost through opt-in to avoid collecting data beyond the socially optimal level.

The third and last potential scheme is a *de-correlation mechanism* (Acemoglu et al. forthcoming; Ichihashi 2020b). The idea is to remove the correlation that is at the root of negative data externalities. Specifically, a trusted third party collects all personal data and then computes transformed variables for each consumer removing the correlation with the information of other consumers and only shares

[23] Fainmesser et al. (2021) suggest that a similar situation to the *personalized Pigovian tax* could be achieved by imposing a tax on online platforms based on the amount of data they collect.

the transformed data. This scheme always improves the equilibrium surplus.

An example of the de-correlation mechanism is an algorithm called the *Randomized Aggregatable Privacy-Preserving Ordinal Response (RAPPOR)* proposed by Erlingsson et al. (2014). RAPPOR is a technology developed in an open-source web browser development project called Chromium, which consists of two privacy-protection mechanisms that satisfy local differential privacy. The reason for the two privacy-protection steps is, intuitively, to avoid *attacks* that exploit the correlation of data. Although it is currently limited to a specific correlation structure, the algorithm can be regarded as an attempt to reduce inefficiency due to negative data externalities.

# 5 Concluding Remarks

This paper presented a survey of the field known as the economics of privacy. The economics of privacy has evolved rapidly in recent years in line with growing global concerns over the handling of personal data on the Internet. The insights can be summarized as follows:

- Privacy means different things to different people.
- The degree of privacy protection can be described by *differential privacy*.
- It is difficult to estimate or observe the degree of privacy protection people desire.
- There is a *privacy paradox*: While people claim to be very concerned about their privacy, they nevertheless undertake very little to protect their personal data.
- Privacy protection imposes costs on firms.
- However, firms may benefit from efforts to protect consumers' privacy.
- Market mechanisms cannot determine the socially optimal level of privacy protection and cannot resolve the problem of *negative data externalities*.
- Determining the socially optimal level of privacy protection is a difficult task.
- Reducing the inefficiency arising from negative data externalities is also a challenge.

## References


**Abowd, John M. and Ian M. Schmutte** (2019) "An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices," *American Economic Review* 109(1), 171–202.

**Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar** (forthcoming) "Too Much Data: Prices and Inefficiencies in Data Markets," *American Economic Journal: Microeconomics*.

**Acquisti, Alessandro** (2004) "Privacy in Electronic Commerce and the Economics of Immediate Gratification," *Proceedings of the 5th ACM Conference on Electronic Commerce*, 21–29.

**Acquisti, Alessandro, Leslie K. John, and George Loewenstein** (2013) "What is Privacy Worth?," *Journal of Legal Studies* 42(2), 249-274.

**Acquisti, Alessandro, Curtis Taylor, and Lian Wagman** (2016) "The Economics of Privacy," *Journal of Economic Literature* 54(2), 442-492.

**Allcott, Hunt and Matthew Gentzkow** (2017) "Social Media and Fake News in the 2016 Election," *Journal of Economic Perspectives* 31(2), 211-36.

**Aridor, Guy, Yeon-Koo Che, and Tobias Salz** (2020) "The Effect of Privacy Regulation on the Data Industry: Empirical Evidence from GDPR," NBER Working Paper 26900.

**Arrieta-Ibarra, Imanol, Leonard Goff, Diego Jiménez-Hernández, Jaron Lanier, and E. Glen Weyl** (2018) "Should We Treat Data as Labor? Moving beyond 'Free'," *AEA Papers and Proceedings* 108, 38-42.

**Athey, Susan, Christian Catalini, and Catherine E. Tucker** (2017) "The Digital Privacy Paradox: Small Money, Small Costs, Small Talk," NBER Working Paper 23488.

**Barnes, Susan B.** (2006) "A Privacy Paradox: Social Networking in the United States," *First Monday* 11(9).

**Becker, Gary. S.** (1980) "Privacy and Malfeasance: A Comment," *Journal of Legal Studies* 9(4), 823–826.

**Bergemann, Dirk, Alessandro Bonatti, and Tan Gan** (2020) "The Economics of Social Data," *arXiv*:2004.03107v1.

**Catalini, Christian and Catherine E. Tucker** (2016) "Seeding the S-Curve? The Role of Early Adopters in Diffusion," NBER Working Paper 22596.

**Catalini, Christian and Catherine E. Tucker** (2017) "When Early Adopters Don't Adopt," *Science* 357(6347), 135-136.

**Chen, Long, Yadong Huang, Shumiao Ouyang, and Wei Xiong** (2021) "The Data Privacy Paradox and Digital Demand," NBER Working Paper 28854.

**Chiou, Lesley and Catherine E. Tucker** (2017) "Search Engines and Data Retention: Implications for Privacy and Antitrust," NBER Working Paper 23815.

**Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim** (2019) "Privacy and Personal Data Collection with Information Externalities," *Journal of Public Economics*, 173, 113-124.

**Duchi, John C., Michael I. Jordan, and Martin J. Wainwright** (2013) "Local Privacy and Statistical Minimax Rates," *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS)*, 429-438.

**Dwork, Cynthia** (2006) "Differential Privacy," *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, 1-12.

**Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith** (2006) "Calibrating Noise to Sensitivity in Private Data Analysis," *Proceedings of the Third Theory of Cryptography Conference TCC, volume 3876 of Lecture Notes in Computer Science*, 265-284.

**Englehardt, Steven and Arvind Narayanan** (2016) "Online Tracking: A 1-million-site Measurement and Analysis," *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 1388–1401.

**Erlingsson, Ulfar, Vasyl Pihur, and Aleksandra Korolova** (2014) "RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response," *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security - CCS'14*, 1054–1067.

**European Central Bank** (2021) "Eurosystem Report on the Public Consultation on a Digital Euro," Available at https://www.ecb.europa.eu/paym/digital_euro/html/pubcon.en.html.

**Fainmesser, Itay Perah and Andrea Galeotti, and Ruslan Momot** (2021) "Digital Privacy," HEC Paris Research Paper No. MOSI-2019-1351. Available at https://ssrn.com/abstract=3459274.

**Financial Stability Board** (2019) "BigTech in Finance," Available at https://www.fsb.org/2019/12/bigtech-in-finance-market-developments-and-potential-financial-stability-implications/.

**Ghosh, Arpita and Aaron Roth** (2011) "Selling Privacy at Auction," *Proceedings of the 12th ACM Conference on Electronic Commerce*, 199–208.

**Goldberg, Samuel, Garrett Johnson, and Scott Shriver** (2019) "Regulating Privacy Online: The Early Impact of the GDPR on European Web Traffic & E-Commerce Outcomes," Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3421731.

**Goldfarb, Avi and Catherine E. Tucker** (2011) "Privacy Regulation and Online Advertising," *Marketing Science* 57(1), 57-71.

**Goldfarb, Avi and Catherine E. Tucker** (2012a) "Shifts in Privacy Concerns," *American Economic Review* 102(3), 349–353.

**Goldfarb, Avi and Catherine E. Tucker** (2012b) "Privacy and Innovation," *Innovation Policy and the Economy* 12, 65-90.

**Heckman, James, J.** (1979) "Sample Selection Bias as a Specification Error," *Econometrica* 47(1), 153–161.

**Heffetz, Ori and Katrina Ligett** (2014) "Privacy and Data-Based Research," *Journal of Economic Perspectives* 28(2), 75-98.

**Hsu, Justin, Marco Gaboardi, Andreas Haeberlen, Sanjeev Khanna, Arjun Narayan, Benjamin C. Pierce, and Aaron Roth** (2014) "Differential Privacy: An Economic Method for Choosing Epsilon," *Proceedings of 27th IEEE Computer Security Foundations Symposium*, 398-410.

**Huberman, Bernardo A., Eytan Adar, and Leslie Fine** (2005) "Valuating Privacy," *IEEE Security and Privacy* 3(5), 22–25.

**Hui, Kai-Lung and Ivan Paak Liang Png** (2006) "The Economics of Privacy," In *Handbooks in Information Systems: Volume 1: Economics and Information Systems*, edited by Terrence Hendershott, 471-498. Bingley, UK: Emerald.

**Ichihashi, Shota** (2020a) "Online Privacy and Information Disclosure by Consumers," *American Economic Review* 110(2), 569-595.

**Ichihashi, Shota** (2020b) "The Economics of Data Externalities," Available at https://shota2.github.io/research/externality.pdf

**Ichihashi, Shota** (2020c) "Dynamic Privacy Choices," Available at https://shota2.github.io/research/dynamicPrivacy.pdf

**Jia, Jian, Ginger Zhe Jin, and Liad Wagman** (forthcoming) "The Short-Run Effects of the General Data Protection Regulation on Technology Venture Investment," *Marketing Science*.

**Jones, Charles I. and Christopher Tonetti** (2020) "Nonrivalry and the Economics of Data," *American Economic Review* 110(9), 2819-2858.

**Kasiviswanathan, Shiva Prasad, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith** (2011) "What Can We Learn Privately?," *SIAM Journal on Computing* 40(3), 793-826.

**Konečný, Jakub, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik** (2016) "Federated Optimazation: Distributed Machine Learning for On-Device Intelligence," *arXiv*:1610.02527.

**Konečný, Jakub, H. Brendan McMahan, Felix X. Yu, Ananda Theertha Suresh, Dave Bacon, and Peter Richtárik** (2017) "Federated Learning: Strategies for Improving Communication Efficiency," *arXiv*:1610.05492v2.

**Kummer, Michael and Patrick Schulte** (2019) "When Private Information Settles the Bill: Money and Privacy in Google's Market for Smartphone Applications," *Management Science* 65(8), 3470-3494.

**Lanier, Jaron** (2013) *Who Owns the Future?*, New York: Simon & Schuster.

**Ligett, Katrina and Aaron Roth** (2012) "Take It or Leave It: Running a Survey When Privacy Comes at a Cost," *Proceedings of the 8th International Conference on Internet and Network Economics*, 378–391.

**Lin, Tesary** (2021) "Valuing Intrinsic and Instrumental Preferences for Privacy," Available at https://tesarylin.github.io/uploads/JMP-Tesary.pdf

**McMahan, H. Brendan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas** (2016) "Federated Learning of Deep Networks using Model Averaging," *arXiv*:1602.05629v1.

**McSherry, Frank and Kunal Talwar** (2007) "Mechanism Design via Differential Privacy," *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, 94-103.

**Miller, Amalia R. and Catherine E. Tucker** (2009) "Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records," *Management Science* 55(7), 1077-1093.

**Miller, Amalia R. and Catherine E. Tucker** (2011) "Can Health Care Information Technology Save Babies?," *Journal of Political Economy* 119(2), 289-324.

**Narayanan, Arvind and Vitaly Shmatikov** (2008) "Robust Deanonymization of Large Sparse Datasets," *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, 111–125.

**Pai, Mallesh M. and Aaron Roth** (2013) "Privacy and Mechanism Design," *ACM SIGecom Exchanges* 12(1), 8-29.

**Posner, Eric A. and E. Glen Weyl** (2018) *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*, Princeton, NJ: Princeton University Press.

**Posner, Richard A.** (1978) "The Right of Privacy," *Georgia Law Review* 12(3), 393-422.

**Posner, Richard A.** (1981) "The Economics of Privacy," *American Economic Review* 71(2), 405-409.

**Ramadorai, Tarun, Antoine Uettwiller, and Ansgar Walther** (2019) "The Market for Data Privacy," CEPR Discussion Paper 13588.

**Rosen, Jeffrey** (2012) "The Right to be Forgotten," *Stanford Law Review Online* 64, 88.

**Stigler, George J.** (1980) "An Introduction to Privacy in Economics and Politics," *Journal of Legal Studies* 9(4), 623–44.

**Tucker, Catherine E.** (2014) "Social Networks, Personalized Advertising, and Privacy Controls," *Journal of Marketing Research* 51(5), 546-562.

**Varian, Hal R.** (1996) "Economic Aspects of Personal Privacy," In *Privacy and Self-Regulation in the Information Age*, Washington, DC: US Department of Commerce, National Telecommunications and Information Administration.

**Wei, Miranda, Madison Stamos, Sophie Veys, Nathan Reitinger, Justin Goodman, Margot Herman, Dorota Filipczuk, Ben Weinshel, and Michelle L. Mazurek** (2020) "What Twitter Knows: Characterizing Ad Targeting Practices, User Perceptions, and Ad Explanations through Users' Own Twitter Data," *Proceedings of the 29th USENIX Security Symposium* 145-162.

**Xiong, Xingxing, Shubo Liu, Dan Li, Zhaohui Cai, and Xiaoguang Niu** (2020) "A Comprehensive Survey on Local Differential Privacy," *Security and Communication Networks*.

# A Local Differential Privacy

Local differential privacy (LDP) is a state-of-the-art privacy protection technique in which each person locally perturbs her/his *raw* personal data with a differential privacy mechanism and transfers the perturbed data to a central server to construct a dataset (Kasiviswanathan et al. 2011; Duchi et al. 2013). In an LDP system, people's raw personal data are stored only on their local devices.[24] The dataset on the central server contains only data after privacy-protection noise has been added, so that the dataset administrator cannot observe the raw personal data. This means that no raw personal data can be leaked from the dataset.

As an example of the implementation of LDP, RAPPOR, which was discussed in Section 4.3, is an algorithm that adds noise in the browser of each local device and sends the processed data to the server. In addition to RAPPOR, recently, software developed by tech companies such as Apple, Microsoft, and Samsung has implemented algorithms that satisfy LDP (Xiong et al. 2020).

[24] The idea of doing some kind of learning or aggregation on a central server while keeping sensitive data on local devices is similar to a technique called *federated learning* proposed by McMahan et al. (2016) and Konečný et al. (2016, 2017). In federated learning, the model to be trained is shared, data in each local device trains separately, and the model parameters are updated by sending the results of training to the central server. In the field of machine learning, federated learning is considered to make major contributions regarding privacy protection.