THE PRIVACY ELASTICITY OF BEHAVIOR: CONCEPTUALIZATION AND APPLICATION

Inbal Dekel
Rachel Cummings
Ori Heffetz
Katrina Ligett

WORKING PAPER 30215

NBER WORKING PAPER SERIES

THE PRIVACY ELASTICITY OF BEHAVIOR: CONCEPTUALIZATION AND APPLICATION

Inbal Dekel
Rachel Cummings
Ori Heffetz
Katrina Ligett

Working Paper 30215
http://www.nber.org/papers/w30215

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2022

Katrina Ligett is currently visiting faculty at Google. We thank participants in HUJI's BEE, JESC, and Rationality Retreat, the Israel Economic Association's Annual Meeting, and FAIR Midway Conference (NHH) for useful feedback and ideas, and Bnaya Dreyfuss, Ofer Glicksohn and Guy Ishai for comments. For financial support, we acknowledge the following. Dekel: The United States Air Force and DARPA under contracts FA8750-16-C-0022 and FA8750-19-2-0222, and the Federmann Cyber Security Center in conjunction with the Israel national cyber directorate. Cummings: NSF grants CNS-1850187 and CNS-1942772 (CAREER), the Defense Advanced Research Projects Agency under contract number W911NF-21-1-0371, a Mozilla Research Grant, a JPMorgan Chase Faculty Research Award, a Google Research Fellowship, and an Apple Privacy-Preserving Machine Learning Award; part of this work was completed while Cummings was affiliated with the California Institute of Technology and the Georgia Institute of Technology, and while visiting the Simons Institute for the Theory of Computing and the Hebrew University of Jerusalem. Heffetz: Israel Science Foundation (grant no. 1680/16), and the S.C. Johnson Graduate School of Management. Ligett: NSF grants CNS-1254169 and CNS-1518941, US-Israel Binational Science Foundation grant 2012348, Israel Science Foundation (ISF) grant #1044/16, Israeli Science Foundation (ISF) grant #2861/20, the United States Air Force and DARPA under contracts FA8750-16-C-0022 and FA8750-19-2-0222, a Google Faculty Research award, Simons Foundation Mathematical and Physical Science Collaboration number 733792, the Federmann Cyber Security Center in conjunction with the Israel national cyber directorate, a gift to the McCourt School of Public Policy and Georgetown University, and a grant from the Israeli National Data Science initiative; part of this work was done while Ligett was visiting the Simons Institute for the Theory of Computing. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Inbal Dekel, Rachel Cummings, Ori Heffetz, and Katrina Ligett. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Privacy Elasticity of Behavior: Conceptualization and Application
Inbal Dekel, Rachel Cummings, Ori Heffetz, and Katrina Ligett
NBER Working Paper No. 30215
July 2022
JEL No. C91,D82,Z00

ABSTRACT

We propose and initiate the study of privacy elasticity—the responsiveness of economic variables to small changes in the level of privacy given to participants in an economic system. Individuals rarely experience either full privacy or a complete lack of privacy; we propose to use differential privacy—a computer-science theory increasingly adopted by industry and government—as a standardized means of quantifying continuous privacy changes. The resulting privacy measure implies a privacy-elasticity notion that is portable and comparable across contexts. We demonstrate the feasibility of this approach by estimating the privacy elasticity of public-good contributions in a lab experiment.

Inbal Dekel
The Bogen Family Department of Economics
The Hebrew University of Jerusalem
Jerusalem
Israel
inbal.dekel1@mail.huji.ac.il

Rachel Cummings
Industrial Engineering and
Operations Research,
Columbia University
535E S.W. Mudd
Mail Code: 4704
New York, NY 10027
rac2239@columbia.edu

Ori Heffetz
S.C. Johnson Graduate School of Management
Cornell University
324 Sage Hall
Ithaca, NY 14853
and The Hebrew University of Jerusalem
and also NBER
oh33@cornell.edu

Katrina Ligett
The Rachel and Selim Benin School of
Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem
Israel
katrina.ligett@mail.huji.ac.il

A data appendix is available at http://www.nber.org/data-appendix/w30215

We increasingly live our lives under constant digital surveillance. Perfect privacy is rarely an option, but neither (for the most part) do our actions appear on the front page of the *New York Times*. The reality is somewhere in between, and our behavior might respond accordingly. This paper is focused on the *privacy elasticity* of behavior: What is the percentage change in a behavioral outcome in response to a one-percent change in privacy?

Elasticity is a fundamental concept in economics, and private versus public behavior has long been studied by economists. However, to the best of our knowledge, the combination—elasticity with respect to privacy, or simply *privacy elasticity*—has been all but absent from economists' vocabulary. The reason may be the lack of a standardized way for economists to think about, conceptualize, and quantify intermediate privacy levels. Indeed, what does a "one-percent change in privacy" even mean?

Our first contribution is to propose an answer to this question, and to derive from it a workable definition of privacy elasticity. Our second contribution is to demonstrate how such privacy elasticity can be empirically estimated.

Mirroring these two contributions, the paper consists of two main sections, followed by a concluding discussion. In Section 1 we conceptualize privacy. We start by importing a continuous, standardized measure of privacy guarantees developed by computer scientists: ϵ-*differential privacy* (Dwork et al., 2006*a*). This measure is being widely adopted, including in recent high-profile deployments at the US Census (Dajani et al., 2017), Apple (Apple Differential Privacy Team, 2014), and Google (Erlingsson, Pihur and Korolova, 2014; Fanti, Pihur and Erlingsson, 2016).

Intuitively, differential privacy protects the privacy of individual data elements by adding noise to any record or publication of either the data itself or statistics based it. This noise is guaranteed to provide a provable upper bound on the ratio between an observer's posterior beliefs and what they would have been if any one data element were actually a completely different value. Differential privacy thus provides a standardized, portable, and readily measurable privacy parameter: the upper bound on this ratio, parametrized as e^{ϵ}, with $\epsilon \geq 0$. When $\epsilon = 0$, the ratio is 1 and privacy is complete. As ϵ increases, the noisy output—and hence the public signal provided by the individual's data—can be increasingly

informative.[1]

After reviewing the definition of differential privacy, in the rest of Section 1 we discuss some basic properties of the notion and, importantly, interpret the meaning of a "one-percent change in privacy." We then derive the implied formal definition of privacy elasticity. We conclude the section with examples of how differential privacy is being currently applied by major tech firms, who already use ϵ to both *quantify* the level of privacy guaranteed to product users and, importantly, to *communicate* that privacy level to the public. In such real-world settings, the notion of privacy elasticity may be readily applied as a useful tool.

In Section 2 we illustrate this applicability, step by step and on a much smaller scale, in a controlled lab environment. We run an experiment where we exogenously vary the privacy parameter, e^{ϵ}, to demonstrate how one might estimate the privacy elasticity of economic behavior in one particular setting. We focus on a public-good game—a setting that has been extensively studied in the lab as an important example of market failure. Importantly, motivated by the idea that making individual contributions public may reduce free riding, the public-good example has been extensively studied in the lab under different *privacy* conditions. We build on past experimental designs that mostly focused on binary private-versus-public conditions. However, armed with a continuous privacy measure from Section 1, we can go beyond past experiments, and estimate the change in contributions resulting from *marginal* privacy changes. We use these estimates to estimate, to the best of our knowledge for the first time, the privacy elasticity of contributions to the public good.

As explained in Section 2, in our experiment ($N = 328$ participants $\times$ 7 rounds $= 2,296$ observations), we exogenously vary both the price of contribution—the amount one has to forgo to generate \$1 in others' takeaway money—and the level of privacy protection, e^{ϵ}, of a public announcement of said contribution. We vary the former between subjects and

[1] Abowd and Schmutte (2019) lament that "our discipline has ceded one of the most important debates of the information age to computer science," and report (p. 174):

> Privacy-preserving data analysis is barely known outside of computer science. A search for "differential privacy" in JSTOR's complete economics collection through December 2017 found five articles. The same query for statistics journals found six. A search of the ACM Digital Library, the repository for the vast majority of refereed conference proceedings in computer science, for the same quoted keyword found 47,100 results.

By basing our proposed definition of privacy elasticity on differential privacy, we hope to also contribute to remedying this situation, and help bring more economists into this important debate.

the latter both between and within subjects. We estimate an average price elasticity of contribution at -0.23 (S.E. $= 0.07$), well within the range of estimates from comparable past experiments. In addition, we estimate an average privacy elasticity of contribution (more precisely, a privacy-loss elasticity of contribution over an arguably plausible range of e^{ϵ}) at 0.07 (S.E. $= 0.01$). This allows us to compare the monetary-contribution response to privacy against the monetary-contribution response to other variables—such as price in our experiment, and income and prices in other studies.

The main insight behind this paper is that the theoretical toolkit of differential privacy can be rather straightforwardly embedded also in empirical economic analysis. This toolkit, which is becoming the industry standard for *protecting* privacy, is also a tool for *quantifying* privacy (or privacy loss), allowing the study of privacy elasticity. In addition, by providing a standardized continuum of formal privacy-protection levels with a natural economic interpretation, this toolkit can be readily imported into the economics lab and—in the future—the field, for studying the behavioral response to changes on the private-public continuum. In our concluding discussion in Section 3 we outline some of the implications of this proposed notion of privacy elasticity.

Finally, this paper may help bring closer several currently mostly disjoint literatures in economics that investigate how privacy can affect behavior. Theoretically, the traditional binary distinction between public and private knowledge (e.g., about an individual's type), which does not readily lend itself to gradual privacy changes, has often been mitigated by introducing various noisy signals. More recently, models of behavior directly integrating privacy considerations—e.g., models of prosocial behavior—introduce a continuous visibility parameter into utility functions. Yet both types of models typically avoid committing to a specific, standardized interpretation of gradual privacy changes, that could be measured and applied across models and contexts.[2]

Empirically, past work in economics that studies changes in behavior under different privacy conditions is, too, mostly focused on either a binary 0/1 privacy notion or an empirical

[2]For example, Bénabou and Tirole's (2006) model introduces a parameter x, which is informally interpreted as measuring "the visibility or salience of [people's] actions: probability that it will be observed by others, number of people who will hear about it, length of time during which the record will be kept, etc." (p. 1656).

continuous privacy measure that is not standardized and is therefore not portable across contexts. In particular, there is a substantial body of experimental findings, but it is mostly from experiments with two extreme conditions: full (or high) privacy versus full (or high) visibility.[3] There is also an observational literature that uses continuous empirical measures of visibility to study a range of economic behaviors.[4] However, as these empirical measures are not based on formal theory, they too are often context-dependent and are not easily linked to either existing theoretical models or other existing empirical work.

1 Privacy Elasticity: Definition and Interpretation

1.1 Differential Privacy

Differential privacy provides a mathematically provable guarantee of privacy protection. This guarantee is typically achieved by systematically adding noise to sensitive data, to computations done on such data, or to the published results of such computations. The guarantee given protects each individual participating in a dataset against inferences made by an observer of the perturbed output.

We briefly introduce differential privacy; see Dwork and Roth (2014) for a textbook treatment, Heffetz and Ligett (2014) for an introduction for empirical researchers, and the current paper's Section 2 below for a concrete, fully worked-out example application. Consider a randomized function M, that is, a function that rather than behaving deterministically, can have output that is drawn from a distribution; that distribution depends on M's input (oth-

[3] In the lab, for instance, in addition to the public-good-game experiments on which we build our own experiment and which we discuss in section 2.1 below, dictator-game participants give less in double-blind trials (Hoffman, McCabe and Smith, 1996) and when given plausible deniability of bad behavior (Dana, Weber and Kuang, 2007; Andreoni and Bernheim, 2009)—and give more when physically facing the recipient (Bohnet and Frey, 1999); and charitable contributions are affected by the coarseness of information (Harbaugh, 1998) and increase by the presence of an audience (Soetevent, 2005) and by contribution visibility (Ariely, Bracha and Meier, 2009). Outside the lab, voter turnout increases when voting records are publicized among family or neighbors (Gerber, Green and Larimer, 2008); enrollment rates for residential energy-conservation programs increase when signers' identities are revealed (Yoeli et al., 2013); and high-school students adhere more to educational-investment norms when choices are revealed to peers (Bursztyn and Jensen, 2015).

[4] Heffetz (2018) reviews eight survey-based visibility measures (of spending by consumers) used in past work to study, e.g., charitable donations and other behaviors. These empirical measures conceptualize visibility as, e.g., the length of time until a behavior (spending) is noticed, or the closeness of interaction needed for it to be noticed.

erwise, M is a trivial function). M takes as input a data element, interpreted as a single individual's data profile, from the domain $\mathcal{X}$ of all possible (realized as well as hypothetical) such profiles. M's randomized output is an element of some range $\mathcal{R}$, interpreted as the published signal about the individual's data profile. M is *ϵ-locally differentially private*[5] (Dwork et al., 2006*a*) if, for any two elements $x, x' \in \mathcal{X}$—that is, for any two conceivable data profiles of an individual—and all possible realizations of the signal $r \in \mathcal{R}$,

$$\frac{\Pr[M(x) = r]}{\Pr[M(x') = r]} \leq e^{\epsilon}.$$

Intuitively, the above definition constrains the function M to produce nearly the same distribution over outputs, no matter what value is input. The extent to which M's output is allowed to depend on its input is controlled by the bound e^{ϵ}, with $\epsilon \geq 0$. Notice that when $\epsilon = 0$, then $e^{\epsilon} = 1$ and the function M must induce identical output distributions no matter what individual data is input, providing perfect privacy, but a perfectly uninformative signal. When $\epsilon = \infty$, M is unconstrained, providing no privacy guarantee, yet allowing a perfectly informative signal. In between, increasingly smaller values of ϵ correspond to stronger privacy guarantees, by making the mechanism's behavior less and less sensitive to the underlying individual data.

As mentioned, the domain $\mathcal{X}$ can be thought of as any sensitive personal data that an individual may not want revealed to, e.g., researchers, the government, Silicon Valley companies, or the public at large. At low e^{ϵ} values, an individual participating in a differentially private computation—function, mechanism, or platform—enjoys a guarantee that nearly the same distribution over revealed outputs would have been induced had her (actual) personal data been replaced with *any other* (hypothetical) data from $\mathcal{X}$. This protective cloak of noise necessarily sacrifices some degree of accuracy of the outputs, but in a manner that is

[5]In other settings, where the goal is to output only aggregate statistics of a database (e.g., the average contribution to the public good in our experiment in Section 2) rather than data that pertains to each individual (e.g., each *participant*'s contribution in our experiment), a variant of this definition can be used where the input to the function M is the entire database, rather than the data of just one individual. The model we consider here is generally known as the *local* model of differential privacy, with this other variant known as the *centralized* model. Importantly, in both models the guarantee is for *differential* privacy: the function M is not restricted in what it could reveal about the world—and hence about individuals—as long as it masks *differences* in any individual's profile.

transparent and quantifiable.

The local differential privacy model gives worst-case guarantees over both all possible data elements x and all possible realizations of the signal r. It may seem unnecessary to protect against what could amount to extremely unlikely events. Indeed, starting with Dwork et al. (2006*b*), a substantial literature relaxes the constraint over signals, allowing for failures of the differential privacy guarantee for extremely unlikely values of r (say, those with probability e^{-32}). On the other hand, relaxing the worst-case guarantee over unlikely values of x would remove privacy protections for exactly those who often need them most—those whose data is unusual.

The differential privacy literature is, intentionally, mostly mute on the issue of how ϵ should be selected—this is viewed as a question for society or for policymakers, not for theoretical computer scientists. However, a tradition has emerged of discussing values of $\epsilon = 0.1$ or 0.2 as "reasonable," and it is common in the literature to prove theorems that only hold for $\epsilon < 1$. In contrast, real-world deployments of differential privacy to date have at times employed much larger ϵ, often by orders of magnitude.[6] This gap between theory and practice highlights the need for research that will help estimate the behavioral impact of changes in e^{ϵ}.[7]

1.2 Privacy Elasticity

In order to discuss *privacy elasticity*—the percent increase in another variable in response to a one-percent change in privacy—we need a privacy metric where small, multiplicative changes are meaningful. We propose using the bound e^{ϵ} on the probability ratio in the differential privacy definition above as this metric.

To interpret a one-percent increase in this bound, consider the following scenario. An

[6]For example, we discuss below a deployment of differential privacy by Apple with an ϵ of 2 *times the number of days a product is in use*, and research revealed that Apple was using an ϵ of 16 per day in another deployment of differential privacy (Tang et al., 2017); test products from the 2018 Census End-to-End Test were released with $\epsilon = 0.25$ (U.S. Census Bureau, 2019), but the final ϵ selected for the 2020 Census's Disclosure Avoidance System was 19.61 (`https://www.census.gov/newsroom/press-releases/2021/2020-census-key-parameters.html`). Our experiment in Section 2 uses privacy conditions roughly corresponding with $\epsilon = 0, 0.5, 1.5, 2.5, 3.5, 5.3, \infty$.

[7]At the same time, this gap may also reflect a high degree of privacy illiteracy among the public, possibly accompanied by little current public sensitivity to—and behavioral impact of—changes in e^{ϵ}, at least in some important real-world settings. We return to this point below.

individual participates in an activity through some platform. Her activity profile x can affect a signal r about her. Example activities include interacting with healthcare providers, taking potentially-tracked online actions such as browsing the web or using a mobile app, responding to a government survey, or contributing to a public good (in the real world or in a lab experiment). The signal could be some message about her that is visible to others, or merely her personal record in some database that she does not control and that someone may access.

The individual takes her participation as given, and chooses an action profile. There are actions that she would prefer to take under absolute privacy protection. However, she is concerned that certain actions, if (and only if) recorded, monitored, tracked, or revealed, might increase the probability of some bad outcome.[8] For example, if her action profile x became known to certain individuals or institutions, she might later be denied medical insurance, face higher prices, be targeted online (legally or malignly), be shamed or merely embarrassed by her sensitive behavior or survey responses, or suffer social repercussions due to being perceived as not sufficiently generous or prosocial.

Consider optimally positive-looking actions: actions that, given the platform's privacy mechanism, minimize the chance of some such bad outcome occurring. Examples include optimally positive-looking browsing behavior, mobile-app use, survey responses, and charitable contributions. Assume a baseline, unavoidable probability q of the bad outcome under such optimally looking actions. Suppose the platform is run with ϵ-differential privacy. Then the individual is guaranteed that no matter what actions she takes, the probability of the bad outcome increases by at most the multiplicative factor e^{ϵ}.[9]

A one-percent increase in privacy loss in this scenario means a one-percent increase in e^{ϵ} used by the platform. This in turn means that the bound on the chance of any output—i.e., a recorded/advertised signal—and thus of any outcome—e.g., being denied insurance, or

[8]More generally, she is concerned that the mere revelation or tracking of certain actions may affect the distribution over future states of the world, *independently* of any direct effects of the same actions taken under a full privacy guarantee.

[9]Formally, $q = \min_{x'} \Pr[\text{bad outcome}|x']$. The differentially private mechanism guarantees that $\forall x$, $\Pr[\text{bad outcome}|x] \leq e^{\epsilon} q$. To see this, recall that $M\colon \mathcal{X} \to \mathcal{R}$ is a probabilistic function from action profiles to signals, and let $F\colon \mathcal{R} \to \mathcal{T}$ be a probabilistic function from signals to outcomes (i.e., to states of the world). If M is differentially private then $\forall x, x'$, for any bad outcome $t \in \mathcal{T}$, observe that $\Pr[F(M(x)) = t] = \sum_{r \in R} \Pr[M(x) = r] \Pr[F(r) = t] \leq \sum_{r \in R} e^{\epsilon} \Pr[M(x') = r] \Pr[F(r) = t] = e^{\epsilon} \Pr[F(M(x')) = t]$.

merely getting funny looks from fellow lab participants—also increases by one percent, from $e^{\epsilon}q$ to $1.01e^{\epsilon}q$. The bound e^{ϵ} is thus a privacy metric where small, multiplicative changes are meaningful.

The implied concept of privacy elasticity has a straightforward, if wordy, interpretation. The elasticity of some variable y with respect to the privacy metric e^{ϵ}, defined as

$$\text{privacy elasticity} = \frac{\partial \log y}{\partial \log e^{\epsilon}} \equiv \frac{\partial \log y}{\partial \epsilon},$$

is the percentage change in y in response to a one-percent change in the upper bound on the ratio between the probability of any outcome induced by the privacy mechanism and what it would have been if an individual's action profile were actually a completely different one.

1.3 Potential Applications

In the next section, we apply a differentially private mechanism in the lab and estimate privacy elasticity as defined above by exogenously varying the mechanism's ϵ. For the economic variable of interest y we use the fraction of a \$10 endowment that lab participants choose to contribute to a public good. Possible outcomes (induced by the privacy mechanism) that a participant might wish to avoid include other participants making negative inferences about her due to an advertised signal that suggests that she made a low contribution, i.e., engaged in free riding. Consistent with past studies, we find that changing the privacy condition from full to no privacy—in our experiment, $\epsilon = 0$ versus $\epsilon = \infty$—causes a sizeable behavioral response in y; going beyond past work, we further find, and estimate, a behavioral responsiveness to intermediate levels of ϵ.

Outside the lab, the extremes of full and no privacy are rarely an option. In the rest of this section we review real-world deployments of differential privacy, and discuss how the notion of privacy elasticity could be applied in those settings.

Differential privacy has rapidly gained traction as an industry-wide standard. For large tech companies, differential privacy can make it possible to obtain insights from data where ethical concerns, internal data-protection procedures, legal restrictions, or reputational considerations might otherwise limit its collection, sharing, or analysis. These are also settings

(e.g., collecting data about inputs typed into phones or computers) where individuals might plausibly modify their behavior or, alternatively, opt out of data sharing, if they felt they were being "watched" without sufficient protection. Hence, the vocabulary of privacy elasticity also helps understand how what can be *learned* from the data might be affected by changes in privacy guarantees.

For example, Apple Watch users have the option to use the ECG app to record their heartbeat and to check the recording for atrial fibrillation (a form of irregular heart rhythm).[10] This data is fed into the Health app on the user's iPhone. Apple might like to know approximately how many Apple Watch users in a particular geographic region are feeding ECG data into the Health app, to help the company understand demand for such health-related features and prioritize new feature deployments. To construct aggregate usage statistics, Apple needs to gather usage information from individual iPhones. However, an individual user might be concerned that by merely using the ECG app they might indicate having a heart condition, which if revealed could potentially lead to adverse treatment by insurers, advertisers, employers, or even potential romantic partners. Apple uses local differential privacy to protect this information before it is gathered, and currently gathers it from users once a day and uses $\epsilon = 2$ per day to protect the identities of the types of health data that a particular user monitors.[11] Since Apple does not wish to know a *specific* user's behavior, but rather aggregate usage patterns, noisy individual data is sufficient.

In this setting, one economic variable of interest y_1 is whether an Apple Watch user is gathering ECG data. If changes in the privacy protection on this information could make users less inclined to use the ECG feature, such changes could have important implications—from making the Apple Watch a less useful product to reducing the incidence of potentially life-saving ECG monitoring by individuals. Different stakeholders, including Apple, its regulators and competitors, law and public-policy makers, and academic researchers might all wish to understand the privacy elasticity of such behavioral changes. Apple, for example, would like to strike a good balance between the information it collects being useful and not harming the appeal of its products. Another economic variable of interest y_2 is whether an

[10] `https://support.apple.com/en-il/HT208955`

[11] `https://developer.apple.com/documentation/healthkit/data_types`
`https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf`

Apple Watch user opts in or out of providing differentially private Health-related data to Apple. If marginal changes in the privacy guarantees on this information might result in a larger fraction of users opting out of sharing data with Apple, this could affect, e.g., Apple's ability to do strategic product planning.

In another example, both Google (Bittau et al., 2017) and Apple have used local differential privacy to protect and gather information about individual user web-browsing behavior. Both companies would like to understand which websites are causing their web browsers to crash so that the relevant bugs can be fixed or the sites can be blocked. However, concerned that visiting certain websites might reveal sensitive or embarrassing information about them, individuals might change their browsing behavior, or their willingness to share browsing data with tech companies, in response to the level of browsing-information privacy guaranteed. Thus, better understanding the privacy elasticity of behavior in these settings could be of interest to different stakeholders.

Additional examples abound, and some of them rely on more sophisticated implementations of differential privacy than we explore in this paper. Windows has used differential privacy to protect information that it collects from millions of Windows 10 devices about users' app usage (Ding, Kulkarni and Yekhanin, 2017) and to protect information that it reveals to managers about how their employees are collaborating (for example, to see what fraction of employees have less than 15 minutes of one-on-one time scheduled with their manager each week).[12] Other major tech companies—from Uber (Johnson et al., 2020) to Snapchat (Pihur et al., 2018) to Salesforce (Sun et al., 2020) to Facebook to Amazon—have built or are seeking to build and deploy tools for differentially private data analysis; and TikTok is posting job ads that describe background in differential privacy as a qualification.[13]

In reality, most users are likely woefully unaware of the level of differential privacy that their sensitive data is accorded and the consequences that this might have, and hence privacy elasticity in practice is likely to be extremely low in many settings. But as privacy literacy rises and the use of differential privacy continues to expand, users will likely learn to adapt

[12] https://blogs.microsoft.com/ai-for-business/differential-privacy/

[13] https://research.facebook.com/blog/2020/06/protecting-privacy-in-facebook-mobility-data-during-the-covid-19-response/

https://www.amazon.science/tag/differential-privacy

https://careers.tiktok.com/position/6995270706842110221/detail, accessed in May, 2022.

their behavior in response to the protections their data receives.

2 Privacy Elasticity in a Public-Good Experiment

2.1 Experimental design

To demonstrate how one may go about measuring privacy elasticity, we embed a differentially private announcement mechanism into an otherwise-standard public-good-game lab experiment. Here we summarize our experimental design. For additional design details, including discussion of why we made certain design decisions, see Appendix A. For full screenshots of the experiment, see Appendix C.

We conducted 41 sessions of the experiment. In each session, a group of eight subjects enters the lab and is seated in front of eight computer stations. Subjects receive identification numbers, and are asked to stand up and introduce themselves by those numbers to all other group members. Subjects then play seven rounds, referred to as "tasks," of a public-good game with their group. In each round, each subject is asked to divide a personal endowment of \$10 between a personal account and a group account, using whole-dollar amounts. Every dollar allocated to the personal account earns one dollar for the subject. Every dollar allocated to the group account earns an *internal return* of \$0.3 for the subject, and an *external return* of either \$0.3 or \$0.5, randomly varied across sessions, for each of the seven other group members.[14] Referring to the amount allocated to the group account as *contribution*, a subject's earning from a round (in \$) is thus:

$$(10 - \text{her contribution}) + 0.3 \times \text{her contribution} + (0.3 \text{ or } 0.5) \times \text{sum of others' contributions.}$$

Hence, when a subject contributes \$1 they end up having paid $(1 - \textit{internal return})$ to generate $(7 \times \textit{external return})$ dollars in others' takeaway money.[15]

To prevent learning and reciprocity, subjects do not receive any feedback between rounds

[14] Varying one return while keeping the other constant is sufficient for estimating the price elasticity of contributions (at the varied price range), while also allowing for estimating the effect of altruism on contributions.

[15] The price of generating \$1 in others' takeaway money is therefore $(1 - \textit{internal return})/(7 \times \textit{external return})$.

(thus the game can be seen as a one-shot game). They are informed in advance that at the end of the experiment, one task (i.e., round) will be randomly chosen that will determine payments for everybody in the session, in addition to receiving a \$10 participation fee.

The differentially private announcement mechanism embedded in the experiment works as follows. When subjects are informed that in the end of the experiment one round will be randomly chosen and will determine payments, they are also informed that public announcements will be made about their selected allocations in the chosen round. Each subject's *announced allocation* may or may not be the same as her actual allocation. In particular, in each round each subject faces a probability of $1 - p$ that her true allocation in that round will be announced, if the round is chosen at the end, and a probability of p that a uniformly randomly selected whole-numbered division of the \$10 will be announced instead. The probability $p \in \{0, 0.05, 0.25, 0.5, 0.75, 0.95, 1\}$ is randomly ordered across session rounds, but is the same in a given round for all subjects in a session.[16]

For clarity, announcements at the end of the experiment use two randomization devices. If $p \in \{0.05, 0.25, 0.5, 0.75, 0.95\}$, each subject first spins a virtual roulette wheel, whose pockets are numbered 1 to 20, to determine whether her selected allocation or a random allocation will be announced. In the latter case, the subject then rolls a virtual 11-sided die numbered 0–10, to determine that random allocation. If $p = 1$, the roulette step is skipped. If $p = 0$, both roulette and die are skipped. Announcements are made by having each subject's *announced allocation* in the chosen round both appear on everyone's screen and read aloud by an experimenter while the subject stands up and faces the other subjects.

In our experiment, the sensitive data x of each individual (which, in this setting, is also her action y), is the number of dollars that she privately chooses to contribute to the public good in a given round. The process of transforming an individual's actual contribution (from the domain $\mathcal{X} = \{\$0, \$1, \ldots, \$10\}$) into the announced noisy signal of her contribution (in the range $\mathcal{R} = \{\$0, \$1, \ldots, \$10\}$) is our differentially private mechanism M. Particularly, the output of this function consists of the individual's true contribution with probability

[16]Note that final payments (to all subjects) are made according to the *true*, rather than the *announced* contributions. Therefore, given subjects' contributions, p affects announcements but not payments. This separation is necessary to avoid confounding preferences for privacy with preferences for money allocations. (Otherwise, selfish decisions would become, e.g., increasingly efficient relative to prosocial decisions as p increased; in the $p = 1$ extreme, *any* contribution would be equally efficient, having no effect on payments.)

$1-p$, and a uniformly randomly selected whole number between 0 and 10 (inclusive) with probability p, for $p \in \{0, 0.05, 0.25, 0.5, 0.75, 0.95, 1\}$. To analyze the level of differential privacy that M guarantees for a particular p, we must consider a pair of possible individual contribution decisions x, x' and an announced contribution r that maximizes $\frac{\Pr[M(x)=r]}{\Pr[M(x')=r]}$. The value $\Pr[M(x) = r]$ is maximized when $x = r$, taking on value $1 - p + p/11$. The value is minimized for any $x' \neq r$, taking on value $p/11$. Thus, the maximum privacy level guaranteed, e^{ϵ}, is $\frac{11-11p+p}{p}$, yielding $\epsilon = \log \frac{11-10p}{p}$. This allows us to translate values of p into values of ϵ for our experiment: $p = 0$ corresponds to $\epsilon = \infty$; $p = 0.05$ to $\epsilon \approx 5.35$; $p = 0.25$ to $\epsilon \approx 3.53$; $p = 0.5$ to $\epsilon \approx 2.49$; $p = 0.75$ to $\epsilon \approx 1.54$; $p = 0.95$ to $\epsilon \approx 0.46$; and $p = 1$ to $\epsilon = 0$.

To ensure that subjects understand the announcements procedure, a simulated announcement is held in each of the first two rounds before subjects make their actual decisions. In each simulated announcement, each subject is randomly assigned a hypothetical allocation (simulating their chosen allocation), and faces the same probability of "true" (simulated) versus uniformly randomized announcement as in the actual task in that round. Subjects then use the roulette wheel and/or die to determine their *simulated announced allocation*, which is then made public, as explained above. In addition, in *all* rounds, right before making allocation decisions, subjects answer a few comprehension questions.

Subjects are told at the beginning of the experiment that they will complete seven tasks, but they do not know that they will be playing seven rounds of the *same* game, and hence do not know that they will face a *range* of probabilities. Their decisions in the first round are therefore independent of the probabilities in the following rounds. We can therefore use the first-round data as between-subjects data, where probabilities are varied only across sessions.

At the end of the experiment, one round is randomly chosen, announcements are made and, while payments are being prepared, subjects complete a brief survey that includes psychological questionnaires assessing personality traits and reputation-, altruism-, and privacy-related preferences. Subjects are then called one by one by their identification number to receive payment in a sealed envelope.

Our experimental design builds on, and extends, several past experiments. First, it

is adapted from Andreoni and Bernheim (2009) to fit a public-good game, rather than a dictator game, as the former enables a higher degree of hiding in the crowd; and to allow privacy guarantees that are independent of subjects' actions.[17] Second, it borrows from Andreoni and Petrie (2004) and Rege and Telle (2004), who manipulate subjects' privacy in a public-good game by either concealing or revealing subjects' contributions, along with their identities, to their group members. Finally, our design follows Goeree, Holt and Laury (2002) in separating the monetary return from contribution to the public good into internal and external returns.

The experiment was programmed with oTree software (Chen, Schonger and Wickens, 2016).

2.2 Experimental Results

We conducted the experiment in the Business Simulation Lab at Cornell University during February 2019. 328 subjects (8 per session $\times$ 41 sessions; average age = 23.1, 65% women) were recruited through an electronic subject-pool system. In total, subjects were asked either 36 or 37 comprehension questions (up to 6 per round), to verify understanding of how payments and announcements work. They had an average of 85.2% correct first-attempt answers over all questions in all rounds. The experiment took up to 90 minutes to complete, and participants earned an average of \$18.1, in addition to a \$10 show-up fee. (Appendix Figure B2 graphs all contributions by all subjects in all rounds.)

Figure 1 displays subjects' average contribution share (out of the \$10 endowment) by privacy condition, pooling across all sessions (i.e., across the two external-return conditions; Appendix Figure B1 recreates the figure by condition). Privacy is measured in the horizontal axis using the ϵ parameter of the differential privacy guarantee. The figure shows that the average share of contribution increases from 30.9% under full privacy ($\epsilon = 0$, labeled "Private" in the figure) to 46.9% under no privacy ($\epsilon = \infty$, labeled "Public").

[17] Andreoni and Bernheim (2009) test audience effects in a dictator game, where with some probability nature intervenes and replaces the dictator's allocation; and the noisy allocations are later announced to all session members. When nature intervenes, it randomizes between two of all the actions available to the dictator. Hence, choosing one of nature's actions gives plausible deniability, while any other action is fully revealing.

Figure 1: Average share of contribution by ϵ

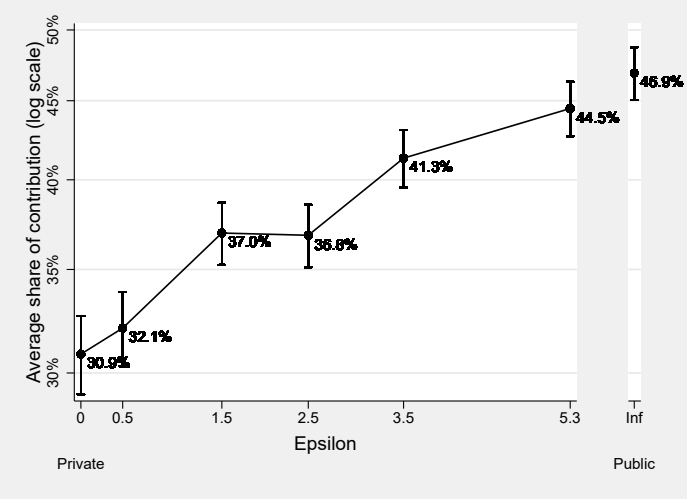

Notes: Capped ranges: ± standard error. $N = 328$ subjects × 7 rounds = 2,296 observations.

Since the average shares of contribution on the y-axis are displayed on a log scale, the slopes represent privacy elasticities as defined in Section 1, i.e., calculated with respect to the probability ratio e^{ϵ}. Elasticities between adjacent privacy levels starting from full privacy ($\epsilon = 0$) are as follows (standard errors in parentheses): 0.08 (0.17), 0.13 (0.07), −0.004 (0.07), 0.11 (0.06), 0.04 (0.03); focusing on the finite extremes of $\epsilon = 0$ and $\epsilon = 5.3$, we estimate an overall average privacy elasticity of contribution at 0.07 (0.01).[18] That the rightmost point in the figure (contribution share at $\epsilon = \infty$) is only 2.5 percentage points above, and not statistically different from, the point immediately to its left (contribution share at $\epsilon = 5.3$) suggests that elasticity quickly drops towards 0 above $\epsilon = 5.3$. (That the rightmost point is so far below 100 percent contribution furthermore suggests that this quick drop is not due to a ceiling effect.) Looking at all slopes, elasticity possibly starts dropping already somewhat below $\epsilon = 5.3$.

Table 1 presents results from OLS regressions where the dependent variable is log(1 + amount contributed). Privacy is measured by ϵ and, aside from the extreme of no privacy ($\epsilon = \infty$, indicated by a dummy variable), enters linearly. Column (1) shows that on average,

[18] We calculate the privacy elasticity between two privacy levels as the difference in log average contribution divided by the difference in ϵ. We calculate (non-clustered) standard errors using the delta method. (Table 1 below reports clustered S.E.'s.) Similarly, we calculate price elasticity = −0.23 (0.07) by dividing the difference in log average contribution at the two price levels by the difference in log price (see footnote 15).

over our finite ϵ's, a one-percent increase in the probability ratio e^{ϵ} entails a 0.07 (S.E. = 0.01) percentage change in contributions. This result is stable and robust across different specifications (Columns (3)–(5)). In comparison, a one-percent increase in the price of contribution (defined as the price of generating \$1 in others' money) entails a −0.18-to-−0.21 (S.E. = 0.13) percentage change in contributions, however very imprecisely estimated (and not statistically significant; Columns (2)–(4)).

Table 1: Privacy and Price Elasticities (Dep. Var.: log(1 + amount contributed))

	Full Sample					First Round
	(1)	(2)	(3)	(4)	(5)	(6)
Privacy: ϵ	0.07		0.07	0.07	0.07	0.06
	(0.01)		(0.01)	(0.01)	(0.01)	(0.03)
$\epsilon = \infty$	0.41		0.41	0.41	0.41	0.44
	(0.05)		(0.05)	(0.05)	(0.05)	(0.13)
log(Price)		−0.18	−0.18	−0.21		−0.09
		(0.13)	(0.13)	(0.13)		(0.15)
Constant	1.09	1.05	0.85	0.29	1.48	1.43
	(0.04)	(0.17)	(0.16)	(0.48)	(0.02)	(0.67)
Psychological measures				Yes		Yes
Demographic controls				Yes		Yes
Individual fixed effects					Yes	
N observations	2,296	2,296	2,296	2,296	2,296	328
N sessions	41	41	41	41	41	41
R^2	0.03	0.00	0.04	0.19	0.73	0.23

OLS regressions. Dependent variable: log(1 + amount contributed). Standard errors in parentheses, clustered at the session level. Column (6) includes only the first round of each session; all other columns include the full sample. Psychological measures: normalized items from the Big Five Personality Traits questionnaire (John and Srivastava, 1999), Brief Fear of Negative Evaluation Scale (Leary, 1983), Compassionate Love For Strangers-Humanity Scale (Sprecher and Fehr, 2005), and Privacy Orientation Scale (Baruh and Cemalcılar, 2014). Demographic controls: age, gender, Hispanic origin or descent, race, education, economic and social attitudes, and political affiliation. Missing demographic data is represented by dummy variables.

Importantly, the privacy elasticity that we observe is not a mere reaction of subjects to *changes* in privacy levels. Column (6) reruns the specification in Column (4) based on only

the first round of each session, during which subjects did not know that they might (and, in fact, would) face other privacy levels. Column (6)'s privacy-elasticity point estimate—a between-subjects estimate—is close, at 0.06 (S.E. = 0.03), to the within-subject estimates in the other columns, however it is estimated much less precisely. (The price elasticity of contribution in the first round is estimated still less precisely; and its point estimate drops.)[19]

Finally, our privacy-elasticity estimate can be put in context. The past few decades have provided several dozen estimates of income and price elasticities of contributions from lab, field, survey, and administrative data (summarized in Appendix Table B1). For example, Goeree, Holt and Laury (2002), whose experimental design we follow, report estimates implying price elasticity = −0.34 (0.10). This and our estimates above suggest that in this experimental paradigm, contributions are similarly affected by a one-percent increase in price and a 3–5 percent increase in privacy. For another example, the range of six income-elasticity estimates from charitable-contribution lab experiments starting with Eckel and Grossman (2003), 0.60–0.99 (0.03–0.17), suggests a similar proportional effect on contributions of a one-percent increase in income in these studies and a 9–14-percent decrease in privacy in our study.

At the same time, existing income- and price-elasticity estimates vary dramatically across contexts and methods. This highlights the need to estimate elasticities—including privacy elasticity—in a variety of settings.

3 Conclusion

With the ever-expanding collection and storage of personal data, privacy considerations and their potential effects on behavior become increasingly important. Differential privacy—which is quickly becoming the consensus, state-of-the-art tool for privacy protection in large data systems—offers a natural tool for quantifying marginal changes in privacy guarantees. It thus enables estimating the privacy elasticity of economic outcomes.

Admittedly, at present, privacy illiteracy appears to be the norm, and privacy preferences

[19]Running the specification in Column (4) of Table 1 separately for each round (see Appendix Table B2) results in fairly similar privacy-elasticity estimates.

in the field appear easily malleable (Acquisti, John and Loewenstein, 2013). It is therefore not implausible that in many important real-world settings, *at present*, the actual privacy elasticity of behavior is rather low. This, however, may reflect current systems and laws more than fundamental individual preferences. As technologies, awareness, and regulation evolve, privacy elasticity may dramatically increase.

Abowd and Schmutte (2019) discuss the tradeoff faced by statistical agencies between protecting respondent privacy—by injecting more noise into published statistics—and publishing accurate statistics—by minimizing said noise. They advocate for work that will help explore optimal privacy-accuracy tradeoffs. We warmly embrace such an agenda. Our work, however, presents an important departure from their model. While the privacy-accuracy tradeoff they highlight varies the value of ϵ holding the underlying data fixed, in our experiment variation in ϵ is the *cause* of changes in individuals' behavior and thus in their private data. In any of the many settings in which privacy concerns might drive selective participation or changes in behavior, the tradeoff faced by policymakers is more complex than selecting ϵ along a fixed tradeoff curve. Unless behavior is perfectly privacy-inelastic both now and, importantly, well into the future, the choice of ϵ could have complex effects on the data gathered, its accuracy and its representativeness.

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.

Acquisti, Alessandro, Leslie K John, and George Loewenstein. 2013. "What is privacy worth?" *The Journal of Legal Studies*, 42(2): 249–274.

Andreoni, James, and B. Douglas Bernheim. 2009. "Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects." *Econometrica*, 77(5): 1607–1636.

Andreoni, James, and Ragan Petrie. 2004. "Public goods experiments without confidentiality: A glimpse into fund-raising." *Journal of Public Economics*, 88(7-8): 1605–1623.

Apple Differential Privacy Team. 2014. "Learning with Privacy at Scale." *Apple Machine Learning Journal*.

Ariely, Dan, Anat Bracha, and Stephan Meier. 2009. "Doing good or doing well? Image motivation and monetary incentives in behaving prosocially." *American Economic Review*, 99(1): 544–555.

Baruh, Lemi, and Zeynep Cemalcılar. 2014. "It is more than personal: Development and validation of a multidimensional privacy orientation scale." *Personality and Individual Differences*, 70: 165–170.

Bénabou, Roland, and Jean Tirole. 2006. "Incentives and prosocial behavior." *American Economic Review*, 96(5): 1652–1678.

Bittau, Andrea, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. 2017. "Prochlo: Strong privacy for analytics in the crowd." *Proceedings of the 26th Symposium on Operating Systems Principles*, 441–459.

Bohnet, Iris, and Bruno S. Frey. 1999. "The sound of silence in prisoner's dilemma and dictator games." *Journal of Economic Behavior & Organization*, 38(1): 43 – 57.

Bursztyn, Leonardo, and Robert Jensen. 2015. "How does peer pressure affect educational investments?" *The Quarterly Journal of Economics*, 130(3): 1329–1367.

Chen, Daniel L., Martin Schonger, and Chris Wickens. 2016. "oTree–An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance*, 9: 88–97.

Dajani, Aref N., Amy D. Lauger, Phyllis E. Singer, Daniel Kifer, Jerome P. Reiter, Ashwin Machanavajjhala, Simson L. Garfinkel, Scot A. Dahl, Matthew Graham, Vishesh Karwa, Hang Kim, PhilipLeclerc, Ian M. Schmutte,

William N. Sexton, Lars Vilhuber, and John M. Abowd. 2017. "The modernization of statistical disclosure limitation at the U.S. Census Bureau." `https://www2.census.gov/cac/sac/meetings/2017-09/statistical-disclosure-limitation.pdf`.

Dana, Jason, Roberto A. Weber, and Jason Xi Kuang. 2007. "Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness." *Economic Theory*, 33(1): 67–80.

Ding, Bolin, Janardhan Kulkarni, and Sergey Yekhanin. 2017. "Collecting telemetry data privately." *Advances in Neural Information Processing Systems*, 30.

Dwork, Cynthia, and Aaron Roth. 2014. "The algorithmic foundations of differential privacy." *Foundations and Trends in Theoretical Computer Science*, 9(3–4): 211–407.

Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006*a*. "Calibrating noise to sensitivity in private data analysis." *Theory of Cryptography Conference*, 265–284.

Dwork, Cynthia, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. 2006*b*. "Our Data, Ourselves: Privacy Via Distributed Noise Generation." *EUROCRYPT*, 4004: 486–503.

Eckel, Catherine C., and Philip J. Grossman. 2003. "Rebate versus matching: does how we subsidize charitable contributions matter?" *Journal of Public Economics*, 87(3-4): 681–701.

Erlingsson, Úlfar, 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*, 1054–1067.

Fanti, Giulia, Vasyl Pihur, and Úlfar Erlingsson. 2016. "Building a rappor with the unknown: Privacy-preserving learning of associations and data dictionaries." *Proceedings on Privacy Enhancing Technologies*, 2016(3): 41–61.

Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social pressure and voter turnout: Evidence from a large-scale field experiment." *American Political Science Review*, 102(1): 33–48.

Goeree, Jacob K., Charles A. Holt, and Susan K. Laury. 2002. "Private costs and public benefits: Unraveling the effects of altruism and noisy behavior." *Journal of Public Economics*, 83(2): 255–276.

Harbaugh, William T. 1998. "The Prestige Motive for Making Charitable Transfers." *American Economic Review*, 88(2): 277–282.

Heffetz, Ori. 2018. "Expenditure visibility and consumer behavior: New evidence." National Bureau of Economic Research 25161. Accessed 13 December 2019.

Heffetz, Ori, and Katrina Ligett. 2014. "Privacy and data-based research." *Journal of Economic Perspectives*, 28(2): 75–98.

Hoffman, Elizabeth, Kevin McCabe, and Vernon L. Smith. 1996. "Social Distance and Other-Regarding Behavior in Dictator Games." *The American Economic Review*, 86(3): 653–660.

John, Oliver P., and Sanjay Srivastava. 1999. "The Big Five trait taxonomy: History, measurement, and theoretical perspectives." In *Handbook of personality: Theory and research*. Vol. 2, , ed. A. Pervin Lawrence and Oliver P. John, 102–138. Guilford.

Johnson, Noah, Joseph P Near, Joseph M Hellerstein, and Dawn Song. 2020. "Chorus: a programming framework for building scalable differential privacy mechanisms." *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, 535–551.

Leary, Mark R. 1983. "A brief version of the Fear of Negative Evaluation Scale." *Personality and Social Psychology Bulletin*, 9(3): 371–375.

Pihur, Vasyl, Aleksandra Korolova, Frederick Liu, Subhash Sankuratripati, Moti Yung, Dachuan Huang, and Ruogu Zeng. 2018. "Differentially-private "draw and discard" machine learning." *arXiv preprint arXiv:1807.04369*.

Rege, Mari, and Kjetil Telle. 2004. "The impact of social approval and framing on cooperation in public good situations." *Journal of Public Economics*, 88(7): 1625–1644.

Soetevent, Adriaan R. 2005. "Anonymity in giving in a natural context—a field experiment in 30 churches." *Journal of Public Economics*, 89(11–12): 2301–2323.

Sprecher, Susan, and Beverley Fehr. 2005. "Compassionate love for close others and humanity." *Journal of Social and Personal Relationships*, 22(5): 629–651.

Sun, Lichao, Yingbo Zhou, Philip S Yu, and Caiming Xiong. 2020. "Differentially private deep learning with smooth sensitivity." *arXiv preprint arXiv:2003.00505*.

Tang, Jun, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang. 2017. "Privacy loss in Apple's implementation of differential privacy on MacOS 10.12." *arXiv preprint arXiv:1709.02753*.

U.S. Census Bureau. 2019. "2020 CENSUS PROGRAM MEMORANDUM SERIES: 2019.13."

Yoeli, Erez, Moshe Hoffman, David G. Rand, and Martin A. Nowak. 2013. "Powering up with indirect reciprocity in a large-scale field experiment (Supplement 2)." *Proceedings of the National Academy of Sciences*, 110: 10424–10429.