The Privacy Elasticity of Behavior: Conceptualization and Application

Inbal Dekel Rachel Cummings Ori Heffetz Katrina Ligett*

December 24, 2022

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.

KEYWORDS: privacy elasticity, differential privacy, privacy guarantees, visibility, economic experiments, public-good game. JEL CLASSIFICATION: C91, D82, Z00

*Dekel: Bogen Family Department of Economics, The Hebrew University of Jerusalem (e-mail: inbal.dekel1@mail.huji.ac.il). Cummings: Department of Industrial Engineering and Operations Research, Department of Computer Science, and Data Science Institute, Columbia University (e-mail: rac2239@columbia.edu). Heffetz: S.C. Johnson Graduate School of Management, Cornell University, Bogen Family Department of Economics and Federmann Rationality Center, The Hebrew University of Jerusalem, and NBER (e-mail: oh33@cornell.edu). Ligett: School of Computer Science and Engineering and Federmann Rationality Center, The Hebrew University of Jerusalem (e-mail: katrina.ligett@mail.huji.ac.il). 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 and Ligett: 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; part of this work was completed while Heffetz was visiting the School for Public and International Affairs (SPIA) at Princeton University. 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, a Google Faculty Research award, Simons Foundation Mathematical and Physical Science Collaboration number 733792, 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.

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 on 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 provide 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. To make this point concrete, we model and analyze a stylized example, where a firm's optimal choice of ϵ to maximize the accuracy of its collected information is a function of the privacy elasticity of its users' aggregate participation in data sharing. Drawing on that analysis, we close the section by highlighting several conceptual questions of implementation and regulation.

In Section 2 we illustrate the hands-on applicability of privacy elasticity, step by step and on a much smaller scale than the above examples, 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

[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.

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 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 readily be 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 relate 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 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 protects each individual in a dataset against inferences made by an observer of the perturbed output.

[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).

[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.

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 Sections 1.4 and 2.1 below for two concrete, fully worked-out example applications. 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 (otherwise, 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

[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.

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 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]

[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.

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 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.

For now, assume that the individual takes her participation as given, and chooses an action profile. (Below, we give examples where participation itself is a possible action choice.) 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 the mechanism with such optimal looking actions. Suppose the platform is run with ϵ-

[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.

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 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 Section 2 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

[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 $\Pr[F(M(x)) = t] = \sum_{r\in R} \Pr[M(x) = r] \Pr[F(r) = t] \leq \sum_{r\in R} e^\epsilon \Pr[M(x') = r] \Pr[F(r) = t] = e^\epsilon \Pr[F(M(x')) = t]$.

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. We close the section with a detailed analysis of a stylized example.

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.

[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

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 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

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

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 their behavior in response to the protections their data receives.

1.4 A Stylized Example

To make these potential real-world applications of privacy elasticity more concrete, we now model and fully analyze the following simplified optimizing-firm scenario. A mobile-device manufacturer has shipped n units. An unknown fraction θ of the units suffer from a rare defect that users cannot detect on their own. The firm would like to estimate θ. It asks each owner to run a diagnostic that perfectly detects her device defect status $x \in \{0, 1\}$ and sends a (possibly noisy) signal $r \in \{0, 1\}$ from her device to the firm. Device ownership and status are given; a user's only action $y \in \{0, 1\}$ is to opt out of or into running the diagnostic. To further simplify, assume that prior to running the diagnostic, each device is equally likely to be defective—so a user's action y cannot depend on (or correlate with) her device status x.

The firm's objective is to maximize the accuracy of its estimator $\widehat{\theta}$, developed below. Accuracy increases with the number of opt-ins. However, wary users, facing what they perceive as unknown future implications of revealing their device status x and, more generally, suspicious of tech firms' use of their private data, are more likely to opt in when guaranteed more privacy. Aware of this, the firm's engineers embed the diagnostic within the following ϵ-locally differentially private mechanism M: the signal r it sends from an opted-in device equals true device status x with probability $1 - p$ and a uniformly drawn $\{0, 1\}$ with probability p. (We show below that M guarantees a privacy level e^{ϵ} that depends on p.) *The*

[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.

firm's optimization problem: choose e^ϵ *to minimize* $\text{Var}(\widehat{\theta})$.

We construct $\text{Var}(\widehat{\theta})$ as a function of e^ϵ in three steps. First, we find the level e^ϵ that M guarantees. To do so, we look for the pair of device statuses x, x' and signal r that maximize $\frac{\Pr[M(x)=r]}{\Pr[M(x')=r]}$. The value $\Pr[M(x) = r]$ is maximized when $x = r$, taking on value $1 - p + p/2 = (2 - p)/2$, and is minimized when $x' \neq r$, taking on value $p/2$. Thus, M guarantees the privacy level $e^\epsilon = (2 - p)/p$. M can therefore be equivalently described as follows: a participating device sends its true status with probability $e^\epsilon/(e^\epsilon + 1)$, and the opposite status with probability $1/(e^\epsilon + 1)$.

Second, we define the privacy elasticity of *aggregate* participation, $\eta \equiv \partial \log N/\partial\epsilon$, where N is the total number of opt-ins (out of the population n). This definition parallels the standard definition of the wage elasticity of (extensive-margin) aggregate labor supply (e.g., Mui and Schoefer, 2021). Borrowing from that literature, we assume that each individual's opt-in decision y follows a simple "reservation privacy level" rule. Aggregate participation is thus the number N (or, equivalently, fraction N/n) of device owners whose reservation privacy level is met by the mechanism M. Given the n owners' reservation privacy levels, the number of opt-ins N is thus a deterministic function of the privacy level e^ϵ that M guarantees.[14]

Third, we construct the estimator $\widehat{\theta}$ and express its variance in terms of e^ϵ. We follow Wang et al. (2017), who analyze the mechanism M described above, known as Binary Randomized Response (BRR). Given N opt-ins, of which $\widehat{N}_1$ are observed with signal $r = 1$, our estimator $\widehat{\theta}$ is[15]

$$\widehat{\theta} = \left(\frac{\widehat{N}_1}{N} - \frac{1}{e^\epsilon + 1}\right) \cdot \frac{e^\epsilon + 1}{e^\epsilon - 1},$$

[14] Formally, each user i has an upper bound $\bar{\epsilon}_i$, above which she is unwilling to participate. If the (atomless) population distribution of upper bounds is given by some CDF G (and PDF g), then the fraction of opt-ins $N/n = 1 - G(\epsilon)$, and the privacy elasticity $\eta = -g(\epsilon)/(1 - G(\epsilon))$.

[15] It follows from Theorem 1 in Wang et al. (2017) that $\widehat{\theta}$ is unbiased, and from their Theorem 2—whose proof we reproduce here—that its variance is given by the expression above. $\text{Var}(\widehat{\theta}) = \text{Var}\left(\left(\frac{\widehat{N}_1}{N} - \frac{1}{e^\epsilon+1}\right) \cdot \frac{e^\epsilon+1}{e^\epsilon-1}\right) = \frac{(e^\epsilon+1)^2}{N^2(e^\epsilon-1)^2}\text{Var}(\widehat{N}_1)$. Since $\widehat{N}_1$ is the sum of N i.i.d. random variables, of which θN are drawn from a Bernoulli distribution with parameter $\frac{e^\epsilon}{e^\epsilon+1}$, and $(1-\theta)N$ are drawn from a Bernoulli distribution with parameter $\frac{1}{e^\epsilon+1}$, its variance is $\text{Var}(\widehat{N}_1) = N\frac{e^\epsilon}{e^\epsilon+1} \cdot \frac{1}{e^\epsilon+1} = N\frac{e^\epsilon}{(e^\epsilon+1)^2}$. Hence, $\text{Var}(\widehat{\theta}) = \frac{e^\epsilon}{N(e^\epsilon-1)^2}$.

with variance

$$\mathrm{Var}(\widehat{\theta}) = \frac{e^{\epsilon}}{N(e^{\epsilon}-1)^2}.$$

The firm thus faces the usual privacy-accuracy tradeoff. A stronger privacy guarantee (lower e^{ϵ}) means noisier signals which, all else equal, would increase the estimator's variance. However, all else is not equal: more privacy means more participation (higher N) which, by itself, reduces variance. Which effect is stronger depends on the privacy elasticity of aggregate participation, η.

By examining the first-order condition $\partial \mathrm{Var}(\widehat{\theta})/\partial\epsilon = 0$, and using $\eta \equiv \partial \log N/\partial\epsilon$, one can show that the optimal (i.e., variance-minimizing) amount of privacy is $e^{\epsilon} = (\eta-1)/(\eta+1)$ for $\eta < -1$, and $e^{\epsilon} = \infty$ (i.e., no privacy) for $\eta \geq -1$.

This stylized example illustrates three points. First, from a firm's point of view, as discussed above, privacy elasticity estimates in the relevant contexts could become key inputs into economic decisions and their analysis.

Second, from a policymaker's point of view, if a firm's only objective were to maximize the accuracy of the data it collects from users, then unless the privacy elasticity of participation were extremely high—in this example, $\eta < -1$, i.e., above unit elasticity—the firm would provide no privacy ($e^{\epsilon} = \infty$). In other words, unless privacy elasticity—which may at present be extremely low in many real-world contexts—dramatically increases (e.g., through awareness, education, or regulation), data-collecting firms may not be inherently incentivized to provide privacy protection.

Indeed, a policymaker could view the firm's optimization problem from its dual-problem perspective. Starting with some socially desirable privacy level e^{ϵ}, privacy elasticity could then be used to quantify the sufficient change in users' aggregate behavior that would fully incentivize the optimizing firm to provide at least e^{ϵ}. If, without regulation (as in this example), sufficiently high elasticity is deemed unrealistic, then a social planner who values privacy may either simply require firms to provide it (as an imposed constraint) or attempt to increase other costs (e.g., reputational or tax-induced) associated with lax privacy.

Finally, from the public's point of view, while a mechanism that sends both false positive and false negative signals about people *by design* may initially sound counterintuitive or

even alarming, in practice said signals may quickly come to be correctly interpreted. In the BRR mechanism in our example, a high level of privacy protection (i.e., e^ϵ close to 1) means that signals are roughly evenly split between 0 and 1 *regardless of underlying status*. (Recall that a participating device sends its true status with probability $e^\epsilon/(e^\epsilon+1)$, and the opposite status with probability $1/(e^\epsilon+1)$.) Thus, in a high-privacy regime, even people without a sophisticated understanding of the mechanism or of probabilities may learn that an individual's own differentially private signal means essentially nothing about her. (Of course, this intuition would not develop in low-privacy settings, where an individual's signal does carry significant meaning.) Indeed, such experience-based learning may contribute to increased privacy elasticity of participation.

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.[16] Referring to the amount allocated to the group account

[16]Varying one return while keeping the other constant is sufficient for estimating the price elasticity

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.[17]

To prevent learning and reciprocity, subjects do not receive any feedback between rounds (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 instructions repeatedly emphasize to subjects that everyone in the room will know their payment only after leaving the experiment. The experiment is double blind: payments are prepared in a different room by another experimenter who neither sees nor is seen by the subjects; that experimenter places payments in sealed envelopes based on identification numbers, and hands them to the experimenter in the lab (who hands them to the subjects before they leave the lab).

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 then 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 $1 - p$ that her true allocation in that round will be announced, if the round is chosen at the end, and a probability 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.[18]

of contributions (at the varied price range), while also allowing for estimating the effect of altruism on contributions.

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

[18] 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.

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 is her action y, i.e., the number of dollars that she chooses to contribute to the public good in a given round. The experiment's differentially private mechanism M transforms an individual's actual contribution (from the domain $\mathcal{X} = \{\$0, \$1, \ldots, \$10\}$) into its announced noisy signal (in the range $\mathcal{R} = \{\$0, \$1, \ldots, \$10\}$). In particular, M is the 11-values case of a mechanism known as Generalized Random Response (GRR), a generalization of the (2-values) Binary Random Response mechanism from Section 1.4: it outputs the individual's true contribution with probability $1 - p$, and a uniformly randomly selected whole number between 0 and 10 (inclusive) with probability p. To analyze the level of differential privacy that M guarantees for a particular p, we again look for a pair of possible individual contribution decisions x, x' and an announced contribution r that maximize $\frac{\Pr[M(x)=r]}{\Pr[M(x')=r]}$. The value $\Pr[M(x) = r]$ is again maximized when $x = r$, in this case taking on value $1 - p + p/11$, and is minimized for any $x' \neq r$, now 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

(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.) While this separation also means that subjects can learn from their final payment something about others' true contributions, that could only happen after they left the lab. (In theory, in extreme cases where everyone in a session contributed \$0 or \$10, payments would fully reveal, after leaving the lab, everybody's true contributions; in practice, such cases never occur in our data.)

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.[19] 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,

[19] 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.

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

Figure 1: Average Share of Contribution by ϵ

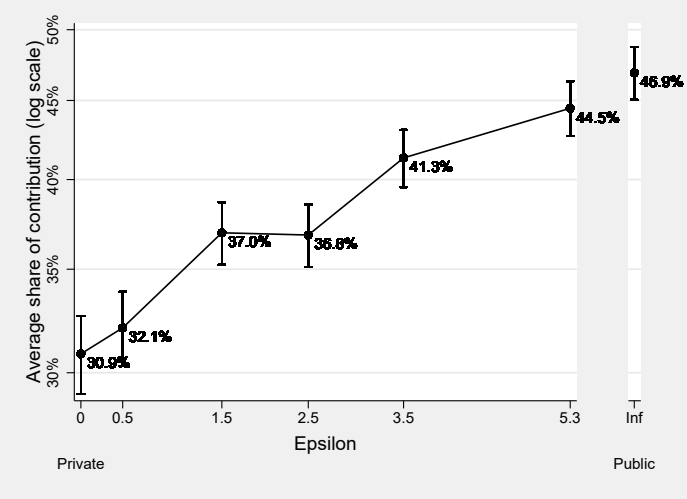

Notes: Capped ranges: ± standard error. $N = 328$ subjects × 7 rounds = 2,296 observations.

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).[20] 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, 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)).

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.)[21]

[20] 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 17).

[21] Running the specification in Column (4) of Table 1 separately for each round (see Appendix Table B1)

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.

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 B2). 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, highlighting the need to estimate elasticities—including privacy elasticity—in a variety of settings. Privacy elasticity may additionally vary with factors ranging from how the private mechanism is implemented and described, to the level of awareness in a society, to cultural differences across societies (see Kachelmeier and Shehata, 1997, for an early cross-society study). Future work may vary these and other factors that are held fixed in our experiment.[22]

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

results in fairly similar privacy-elasticity estimates.

[22] Appendix Figure B1 panels (b)–(r) recreate Figure 1 by the demographics, attitudes, and psychological measures of participants in our experiment. With the exception of the race variable, the figure does not suggest large differences across sample splits.

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. The privacy-accuracy tradeoff they highlight varies the value of ϵ holding the underlying data fixed. In contrast, in our experiment—as in some of the real-world examples we review—variation in ϵ is the *cause* of changes in individuals' behavior and thus in individuals' underlying private data; and in our stylized optimizing-firm example—as in other real-world examples we review—variation in ϵ affects individuals' self-selection into participation in the collected dataset. In any of the many settings in which privacy concerns might drive changes in behavior, selective participation, or both, the tradeoffs faced by policymakers are far more complex than selecting ϵ along a single 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 *underlying* data, as well as on the gathered data, 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 Working Paper 25161.

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.

Kachelmeier, Steven J., and Mohamed Shehata. 1997. "Internal auditing and voluntary cooperation in firms: A cross-cultural experiment." *Accounting Review*, 72(3): 407–431.

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

Mui, Preston, and Benjamin Schoefer. 2021. "Reservation Raises: The Aggregate Labor Supply Curve at the Extensive Margin." National Bureau of Economic Research Working Paper 28770.

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."

Wang, Tianhao, Jeremiah Blocki, Ninghui Li, and Somesh Jha. 2017. "Locally differentially private protocols for frequency estimation." *26th USENIX Security Symposium (USENIX Security 17)*, 729–745.

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.

Appendix: The Privacy Elasticity of Behavior: Conceptualization and Application

Inbal Dekel Rachel Cummings Ori Heffetz Katrina Ligett

A Details of the Experimental Design

Public-good game:

A public-good game is a standard experimental setting in which we can create sensitive data from behavior in the lab. This setting seems appealing, as exposing real sensitive data from everyday life may be unethical, and randomly generating synthetic personal data would give subjects no reason to value privacy. Moreover, as long as true allocations are not revealed, a public-good game allows subjects to hide in the crowd to a greater extent than does, e.g., a dictator game (in which the receiver would always know the amount contributed). Thus, we let subjects play a public-good game where a noisy version of their allocation is announced to the other subjects.

Tasks:

The game consists of seven tasks that differ in the noise parameter of the announcement. The seven noise parameters were chosen to reflect a wide variety of privacy guarantees, ranging from full privacy to no privacy.[1] Considering intermediate noise parameters complements the economic literature, which has focused primarily on the extremes.

The order of tasks in each session is randomly determined. To prevent learning and reciprocity, subjects do not receive any feedback between tasks (thus the game can be seen as a one-shot game). At the end of the game, one task is randomly chosen to determine payments and announcements. This makes it worthwhile for subjects to complete each task as though it will actually be chosen, while allowing us to increase the possible payoffs in each

[1] For more information on the noise parameters, see the subsection titled "Announcements."

task.

Details of the environment:

Subjects play all tasks with the same group, consisting of all 7 other subjects in their session. All group members play in front of computers located in the same computer lab. Announcements are made at the end of the experiment by displaying a noisy version of each subject's allocation decision in the chosen task on everyone's screen; we call this the *announced allocation*. Additionally, an experimenter reads each subject's announced allocation aloud while this subject stands up and faces the other subjects.

These experimental details draw from two experimental studies that manipulate subjects' privacy in a public-good game. In both of these studies, subjects' identities along with their contribution amounts, are either revealed to their group members or not. In the first study, conducted by Rege and Telle (2004), subjects play a one-shot game with a group consisting of all nine other subjects in their session, who are all seated in the same room. Subjects' identification is carried out by asking each subject to come forward, and in front of everyone else, to count the money she contributed and write that amount on a blackboard. In the second study, conducted by Andreoni and Petrie (2004), each subject plays 40 rounds with a group of five subjects, whose composition changes after every eight rounds. All 20 subjects in a session are seated in the same computer lab. Subjects' identification is carried out by displaying their photos and contribution amounts on the screens of all of their group members at the end of each round.

Presumably, having to face your group members while an announcement is made about your allocation is more embarrassing than having your photo displayed on their screens (especially if you have made a low contribution and there is a high chance of announcing your selected allocation). This assumption seems to be supported by the data, as the effect of identification on contribution found by Rege and Telle (2004) is larger than that found by Andreoni and Petrie (2004).

For this reason we chose to ask subjects to stand up and face other subjects while an announcement is made about their allocation.[2] To facilitate this we invite all subjects in a

[2]Subjects are also asked to stand up and introduce themselves by their identification numbers at the

session to a lab, where they can see each other, and assign them to the same group, so that the announcements will only be made in front of other group members (in addition to the experimenter). Furthermore, we chose to have subjects participate on computers as it makes it easier to keep records of subjects' decisions, to check comprehension in real time, and to determine the noisy announced allocations.

Announcements:

Announcements are determined as follows:

In each task of each session, each subject faces a probability $(1-p) \in \{0, 0.05, 0.25, 0.5, 0.75, 0.95, 1\}$ that her selected allocation in this task will be announced (in case this task is chosen at the end of the experiment), and a probability p that a uniformly randomly selected whole-numbered division of the \$10 will be announced instead. The probability p changes from task to task, and it is the same for all subjects in the session.

To promote subjects' understanding of the probabilities and of randomness, a virtual roulette wheel and a virtual die are used as randomization devices. Thus, subjects' announced allocations are determined as follows:

- **Given $p = 0$:**

  Each subject's selected allocation is announced.

- **Given $p = 1$:**

  Each subject is asked to roll a virtual 11-sided die numbered 0-10. The result of this die roll is the subject's announced allocation to the group account.

- **Given $p \in \{0.05, 0.25, 0.5, 0.75, 0.95\}$:**

  Each subject is asked to spin a virtual roulette wheel, whose pockets are numbered from 1 to 20. If the spin result is less than or equal to $(1-p) \cdot 20$ then the subject's selected allocation is announced. Otherwise, a random allocation is announced,[3] in

beginning of the experiment.

[3]To further promote understanding of the probabilities, two rows of circled integers are displayed on subjects' screens alongside the roulette. The first row, that relates to the probability of announcing a subject's selected allocation, contains blue circled numbers that go from 1 to $(1-p) \cdot 20$. The second row, that relates to the probability of announcing a random allocation, contains red circled numbers that go from $(1-p) \cdot 20 + 1$ to 20. The style of these circled numbers matches that of the numbers on the roulette.

which case the subject is asked to roll a virtual 11-sided die numbered 0-10. The result of this die roll is the subject's announced allocation to the group account.

Simulated announcements:

In order to allow subjects to gain experience with the randomization devices (i.e., the roulette and the die) and with the announcement procedure, there is a simulated announcement in each of the first two tasks before subjects make their actual decisions in those tasks. Having two simulated announcements increases the likelihood that subjects will get experience with both the roulette and the die, while keeping the experiment from being too long. In each simulated announcement, a random division of the \$10 is selected for each subject. Each subject is asked to imagine that this division is her selected allocation in the simulated announcement. Each subject faces the same probability p as in the current task. Subjects are then asked to follow through the procedure depicted above to determine their simulated announced allocation (i.e., spin a roulette and/or roll a die). After everyone's simulated announced allocation has been determined, all of the announced allocations are displayed on everyone's screen, and subjects are asked to stand up one at a time while an experimenter reads their simulated announced allocation aloud. Hypothetical allocations in the simulated announcements have no effect on subjects' actual earnings, and this is emphasized to subjects.

Internal and external returns:

To estimate the price elasticity of contributions, while also allowing for a clean estimation of the effect of altruism on contributions, we follow Goeree, Holt and Laury (2002) and slightly modify the standard setup of a public-good game. In a standard public-good game, the monetary return from contribution is the same for the contributor and for all other group members. In such a setup, a change in the common return has two effects, as it changes the net cost of contributing and the monetary benefit to others at the same time.

To avoid this confound, our game separates the monetary return into an 'internal return' for the contributor and a possibly different 'external return' for all other group members. A change in the external return changes only the monetary benefit to others, without affecting the net cost of contributing (and vice versa for a change in the internal return). As it is

enough to vary one of the returns while keeping the other one constant, we chose to keep the internal return constant at 0.3, and to randomly change the external return from session to session so that it would either be 0.3 or 0.5. In each session, the external return is the same for all subjects across all tasks. Given a group size of 8, these returns retain the basic social dilemma structure of the standard public-good game, since the following hold:

(a) The monetary worth of a dollar kept (which is \$1) is greater than the individual's internal return from a dollar contributed. Thus, the dominant strategy for a selfish participant given full privacy is to contribute nothing.

(b) The total return to group members from a dollar contributed $\big($which is: \$$(internal\ return + (8-1)\cdot external\ return)\big)$ is greater than the monetary worth of a dollar kept. Thus, full contribution by all maximizes group earnings.

Instructions:

Instead of providing subjects with instructions regarding all tasks at the beginning of the experiment, we provide them with instructions regarding each task separately at the beginning of that task. In addition, we give subjects a brief introduction at the beginning of the experiment, and also a short explanation at the end about the chosen task and the way announcements are determined.

We give subjects separate instructions regarding each task for a few reasons. First, it helps to simplify the instructions and to promote understanding. Second, it allows us to highlight the probability in each task before decisions are made, and thus ensure that subjects indeed pay attention to the probabilities. Third, it makes subjects' decisions in the first task independent of the probabilities in the other tasks. That is, it prevents subjects from adjusting their allocations in the first task, thinking that they should respond differently to different probabilities. Thus, it allows us to focus on subjects' allocations in the first task as in a between-subjects design. Moreover, comparing subjects' decisions in the first task to their decisions in all other tasks enables us to examine whether there has been some degree of learning, even though subjects do not receive any feedback between tasks.

Instructions are based on a few sources. First, they are adapted from Andreoni and Bernheim (2009) to suit a public-good game (rather than a dictator game), suit the way in

which announcements are determined and their meaning, and to suit having separate and shorter instructions for each task. The second source is Goeree, Holt and Laury (2002), especially the explanations of how payments are determined. The third source is Rege and Telle (2004), especially stressing to subjects that they would maximize their own payment by not contributing but that the group as a whole would benefit from contributions, and their first four examples that further emphasize this point. The final source is Andreoni and Petrie (2004), especially the introduction and decision screens.

Comprehension check:

In addition to providing subjects with a separate set of instructions in each task, a separate comprehension check is conducted in each task right before subjects make their allocation decisions. Conducting a separate comprehension check in each task enables us to make sure that subjects pay attention to the probability of announcing each subject's true allocation (in case this task is chosen at the end of the experiment) and its meaning. Each comprehension check consists of up to six different comprehension questions. Each subject is allowed three attempts to answer each comprehension question in each round before feedback with the correct answer appears on the screen.

The first two questions are inspired by Goeree, Holt and Laury (2002) and they are designed to ensure that subjects understand how payments are calculated. These questions only appear in the first task. The next two questions are designed to ensure that subjects understand what the roulette and die results mean, and more generally that they understand the content of the announcements. Question 3 only appears in tasks in which p is not 0 or 1, and Question 4 only appears in tasks in which p is not 1. Question 5 is designed to ensure that subjects pay attention to the probability in the task. This question does not appear in the first task if the probability in that task is either 0 or 1. Question 6 is designed to ensure that subjects understand how payments are determined. This question only appears in tasks in which p is not 1.[4]

Survey:

[4]For examples of the comprehension questions and possible feedback, see screenshots on pages 27–37.

At the end of the experiment, subjects are asked to answer a brief survey that consists of some standard psychological questionnaires and a number of demographic and attitudinal questions, as well as questions about their reasoning during the experiment. The first questionnaire that subjects are asked to answer is the "Big Five" personality traits questionnaire (John and Srivastava, 1999).[5] The second questionnaire is the Brief Fear of Negative Evaluation (BFNE) Scale, which was found in the literature to be highly correlated with the full-length Fear of Negative Evaluation Scale (Leary, 1983).[6] The third questionnaire is the Compassionate Love For Strangers-Humanity (CLSH) Scale (Sprecher and Fehr, 2005).[7] The fourth questionnaire is the Privacy Orientation Scale (Baruh and Cemalcılar, 2014).[8] Subjects are also asked about their gender, origin, year born, education level, and major, as well as their economic, social, and political attitudes, their comments about the experiment, the way they decided to allocate the money, and what they think the experiment is about.

[5]For the Big Five personality questionnaire, see screenshot on page 44. Items in this questionnaire are rated from 1 (disagree strongly) to 5 (agree strongly). Then, personality traits scores are calculated by the following formulas:
Extroversion: Q1 + (6-Q6) + Q11 + Q16 + (6-Q21) + Q26 + (6-Q31) + Q36;
Agreeableness: (6-Q2) + Q7 + (6-Q12) + Q17 + Q22 + (6-Q27) + Q32 + (6-Q37) + Q42;
Conscientiousness: Q3 + (6-Q8) + Q13 + (6-Q18) + (6-Q23) + Q28 + Q33 + Q38 + (6-Q43);
Neuroticism: Q4 + (6-Q9) +Q14 + Q19 + (6-Q24) + Q29 + (6-Q34) + Q39;
Openness: Q5 + Q10 + Q15 + Q20 + Q25 + Q30 + (6-Q35) + Q40 + (6-Q41) + Q44.

[6]For the BFNE questionnaire, see screenshot on page 46. Items in this questionnaire are rated from 1 (not at all characteristic of me) to 5 (extremely characteristic of me). Then, the score is calculated by the following formula:
Q1 + (6-Q2) + Q3 + (6-Q4) + Q5 + Q6 + (6-Q7) + Q8 + Q9 + (6-Q10) + Q11 + Q12.

[7]For the CLSH questionnaire, see screenshot on page 47. Items in this questionnaire are rated from 1 (not at all true of me) to 7 (very true of me). Then, an average score is calculated for all 21 items.

[8]For the Privacy Orientation questionnaire, see screenshot on page 49. Items in this questionnaire are rated from 1 (strongly disagree) to 5 (strongly agree). Then, scores on privacy dimensions are calculated by the following formulas:
Privacy as a Right: Q1 + Q2 + Q3;
Concern about Own Privacy: Q4 + Q5 + Q6 + Q7;
Other-Contingent Privacy: Q8 + Q9 + Q10 + Q11;
Concern about Others Privacy: Q12 + Q13 + Q14 + Q15 + Q16.

B Additional Tables and Figures

Table B1: Privacy and Price Elasticities (Dep. Var.: log(1 + amount contributed)), by Round

	All	First	Second	Third	Fourth	Fifth	Sixth	Seventh
Privacy: ϵ	0.07	0.06	0.04	0.09	0.10	0.07	0.07	0.09
	(0.01)	(0.03)	(0.02)	(0.03)	(0.04)	(0.03)	(0.02)	(0.03)
$\epsilon = \infty$	0.41	0.44	0.16	0.47	0.38	0.41	0.44	0.62
	(0.05)	(0.13)	(0.22)	(0.15)	(0.14)	(0.14)	(0.15)	(0.12)
log(Price)	−0.21	−0.09	−0.31	−0.27	−0.23	−0.26	−0.12	−0.17
	(0.13)	(0.15)	(0.16)	(0.18)	(0.21)	(0.17)	(0.18)	(0.19)
Constant	0.29	1.43	−0.58	−0.47	0.10	0.45	0.34	0.99
	(0.48)	(0.67)	(0.55)	(0.61)	(0.52)	(0.83)	(0.53)	(0.73)
Psychological measures	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N observations	2,296	328	328	328	328	328	328	328
N sessions	41	41	41	41	41	41	41	41
R^2	0.19	0.23	0.21	0.21	0.28	0.23	0.23	0.21

OLS regressions. Dependent variable: log(amount contributed + 1). Standard errors in parentheses, clustered at the session level. 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.

Table B2: Elasticity Estimates

Study	N	Type	Description	Elasticity			
				Rebate	Matching	Income	Privacy
Peloza and Steel (2005)	1,418,212 (69 studies)	Tax-filer/survey data	Review article[a]	−1.44 (S.D.= 1.21)[b] −1.11[c]			
Goeree, Holt and Laury (2002)	320	Lab	Public-good contributions	−0.34 (0.10)[d,e]			
Eckel and Grossman (2003)	2,016	Lab	Charitable contributions	−0.34 (0.19)	−1.07 (0.18)	0.82 (0.07)	
Eckel and Grossman (2006)	1,080	Lab	Charitable contributions	−1.49 (0.24)	−3.17 (0.24)	0.99 (0.17)	
Karlan and List (2007)	50,083	Natural field	Charitable contributions		−0.23		
Eckel and Grossman (2008)	7,195	Natural field	Charitable contributions	−0.11 (0.04)	−1.05 (0.04)	0.03 (0.01)	
Huck and Rasul (2011)	443	Natural field	Charitable contributions		−0.53 (0.39) to −1.12 (0.44)		
Meer (2014)	371,701	Administrative	Crowdfunfing contributions	−0.78 (0.09)			
Scharf and Smith (2015)	1,737	Hypothetical scenario	Charitable contributions	−0.31 (0.05)	−1.20 (0.09)		
Eckel and Grossman (2017)	1,207	Field	Charitable contributions	−5.12 (0.43)	−5.43 (0.32)	0.19 (0.05)	
Gandullia and Lezzi (2018)	1,456	Lab (online)	Charitable contributions	−0.22 (0.03)		0.60 (0.05)	
	1,208	Lab (online)	Charitable contributions		−1.14 (0.05)	0.80 (0.08)	
Gandullia (2019)	3,568	Lab (online)	Charitable contributions	−0.17 (0.01)		0.60 (0.03)	
	3,480	Lab (online)	Charitable contributions		−1.15 (0.03)	0.77 (0.04)	
This paper	2,296	Lab	Public-good contributions	−0.23 (0.07)[e]			0.07 (0.01)

Standard errors in parantheses unless otherwise stated. The literature distinguishes between prices arising from equivalent rebate and matching subsidies. A rebate rate b is equivalent to a matching rate $m = \frac{b}{(1-b)}$.

[a] Response to changes in tax deductability of charitable contributions.

[b] Weighted mean across all studies.

[c] Weighted mean once outliers are removed.

[d] This elasticity is based on our own calculations and is not reported by the authors.

[e] We define price in Goeree, Holt and Laury's (2002) and our data as follows: $(1 - \textit{internal return})/((N - 1) \times \textit{external return})$, where N is the group size. We then calculate price elasticities as the difference in log contributions at the two price extremes, divided by the difference in log prices.

Figure B1: Variants of Figure 1 by External Return, Demographics, and Psych. Measures

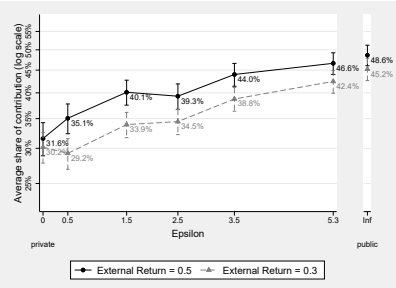

(a) External Return

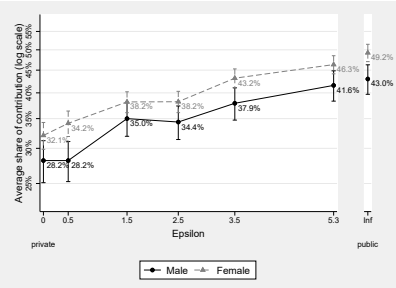

(b) Gender

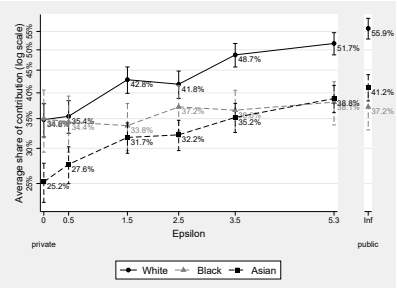

(c) Race

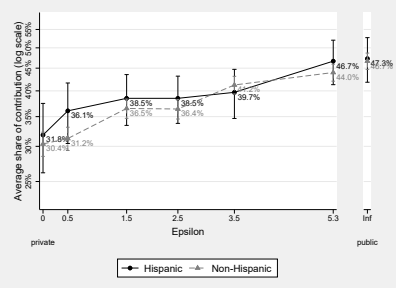

(d) Hispanic

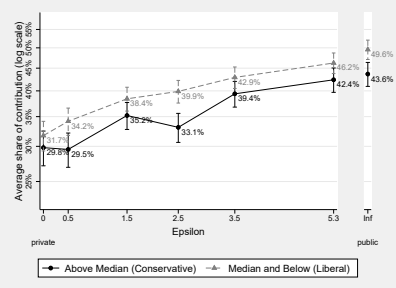

(e) Economic Attitudes

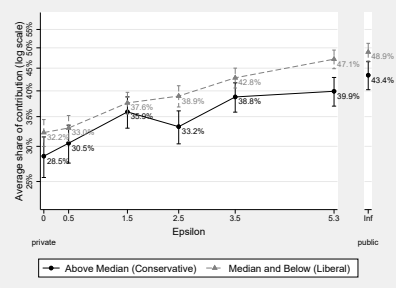

(f) Social Attitudes

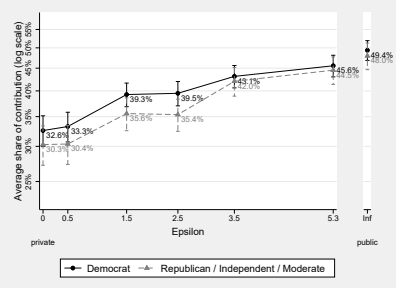

(g) Political Affiliation

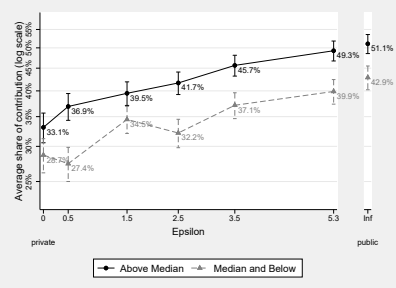

(h) Brief Fear of Negative Evaluation

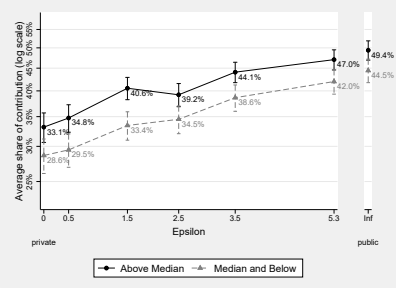

(i) Compassionate Love for Strangers-Humanity

Notes: Capped ranges: ± standard error. Number of participants in each category, by panel (unless split by median): (a) 168 External Return = 0.3, 160 External Return = 0.5; (b) 114 Male, 212 Female (2 Other dropped); (c) 131 White, 32 Black, 138 Asian (1 Native American and 21 Other dropped); (d) 33 Hispanic, 290 non-Hispanic (5 missing responses dropped); (g) 160 Democrat, 20 Republican, 61 Independent, 27 Moderate (8 Other and 51 "None of the above" dropped).

Figure B1: Variants of Figure 1 by External Return, Demographics, and Psych. Measures – Cont.

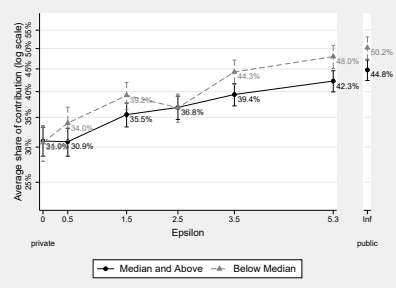

(j) Privacy as a Right

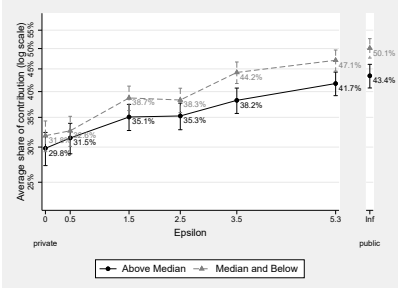

(k) Concern About Own privacy

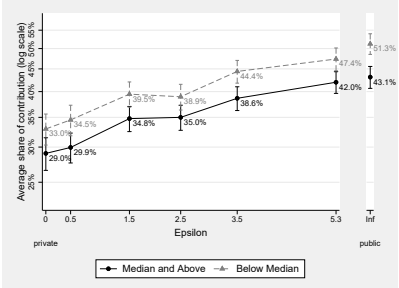

(l) Other-Contingent Privacy

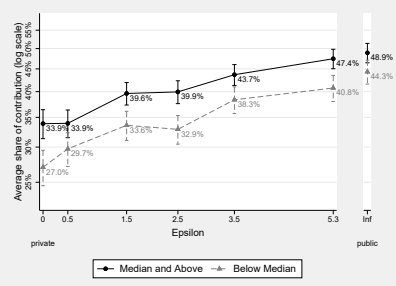

(m) Concern About Others' Privacy

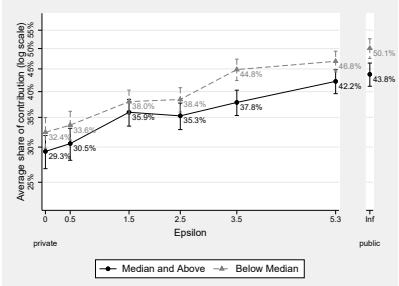

(n) Extroversion

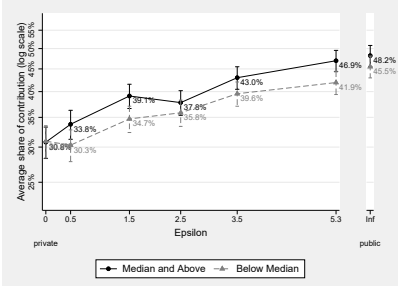

(o) Agreeableness

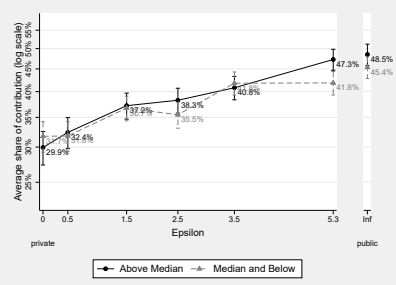

(p) Conscientiousness

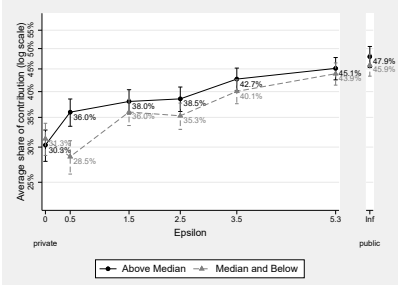

(q) Neuroticism

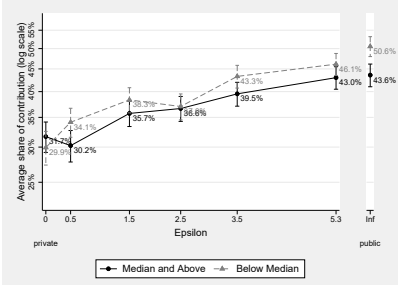

(r) Openness

Notes: Capped ranges: ± standard error.

Figure B2: Individuals' Contributions by ϵ

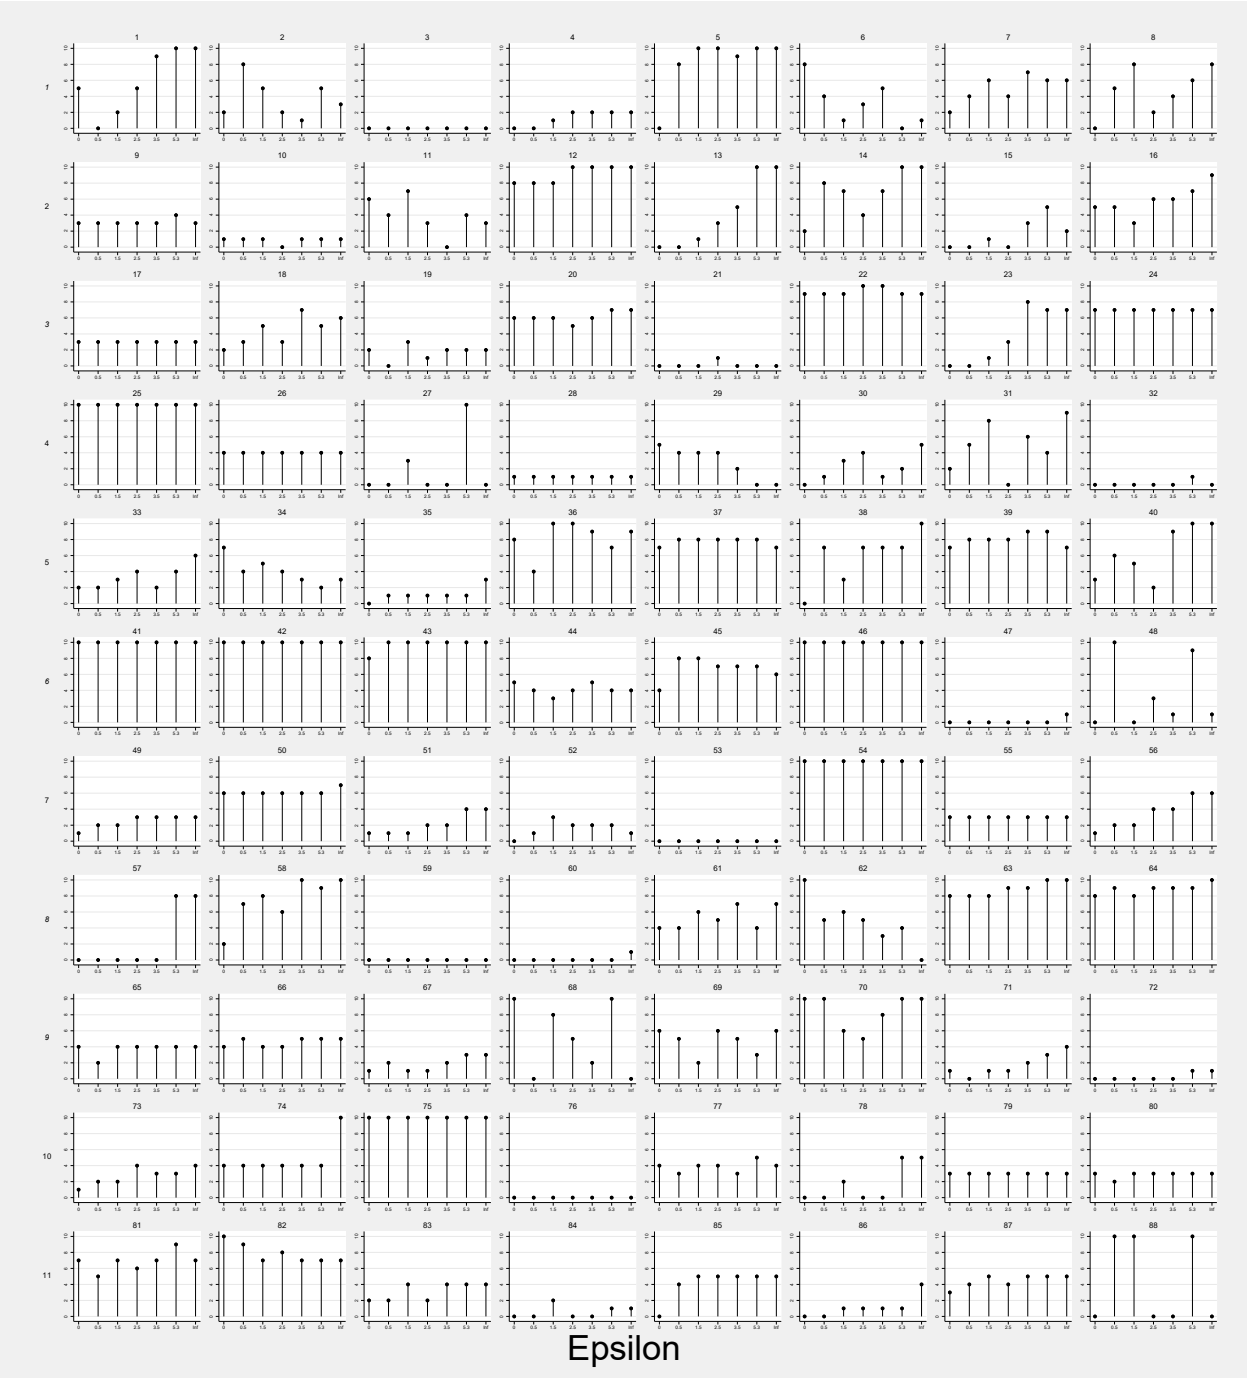

(a) Sessions 1–11

Notes: Each mini-graph represents a single respondent's seven contribution amounts, corresponding with the seven privacy conditions. Respondent number is indicated at the top of the mini-graph. Each row of graphs corresponds with a single session. Session number is indicated at the left of each row, with italicized font indicating a high-external-return session.

Figure B2: Individuals' Contributions by ϵ – Cont.

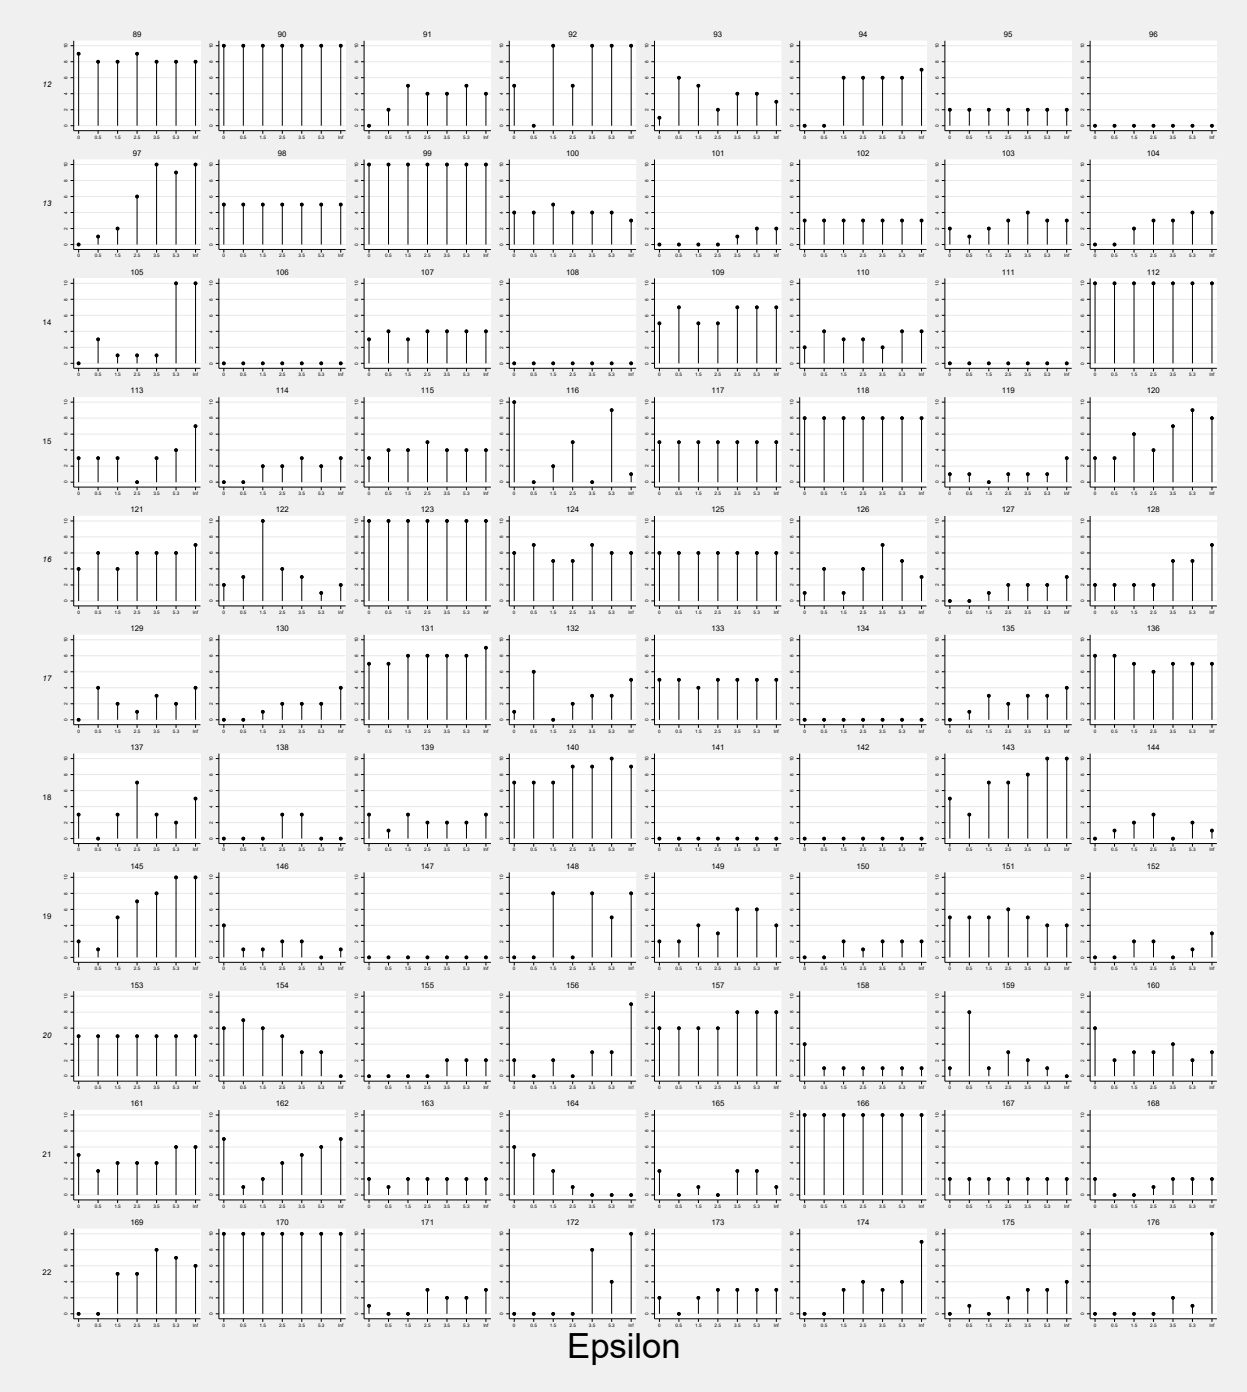

(b) Sessions 12–22

Notes: Each mini-graph represents a single respondent's seven contribution amounts, corresponding with the seven privacy conditions. Respondent number is indicated at the top of the mini-graph. Each row of graphs corresponds with a single session. Session number is indicated at the left of each row, with italicized font indicating a high-external-return session.

Figure B2: Individuals' Contributions by ϵ – Cont.

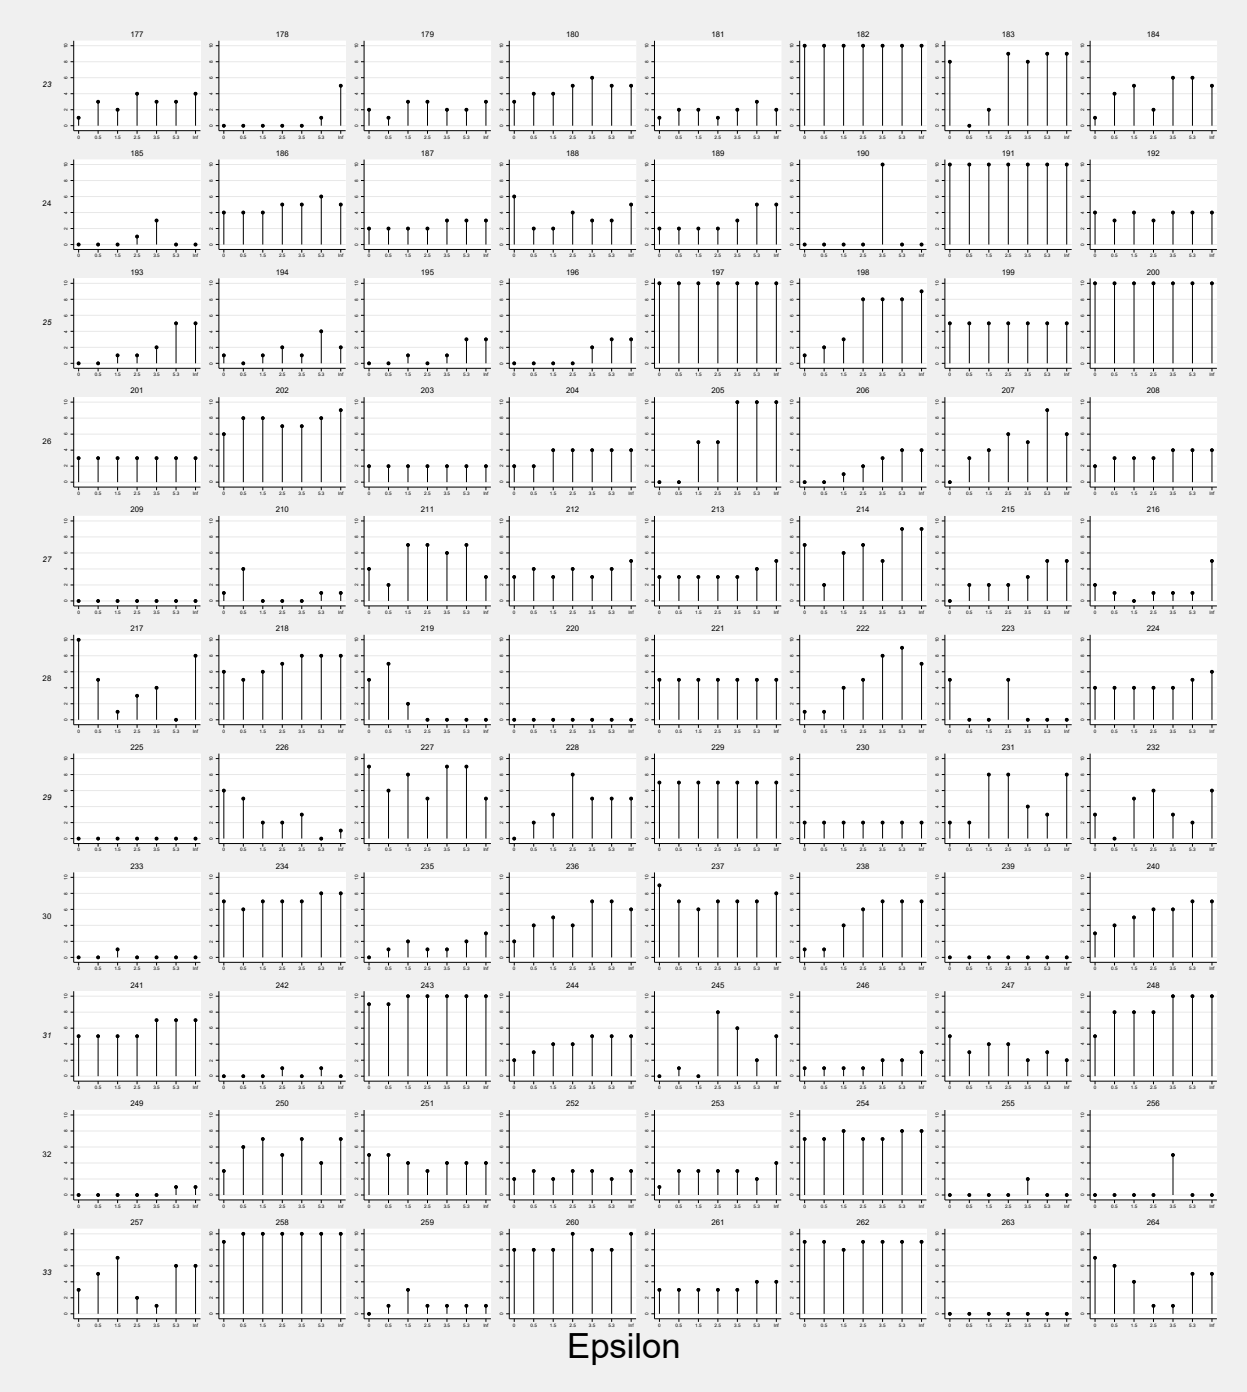

(c) Sessions 23–33

Notes: Each mini-graph represents a single respondent's seven contribution amounts, corresponding with the seven privacy conditions. Respondent number is indicated at the top of the mini-graph. Each row of graphs corresponds with a single session. Session number is indicated at the left of each row, with italicized font indicating a high-external-return session.

Figure B2: Individuals' Contributions by ϵ – Cont.

(d) Sessions 34–41

Notes: Each mini-graph represents a single respondent's seven contribution amounts, corresponding with the seven privacy conditions. Respondent number is indicated at the top of the mini-graph. Each row of graphs corresponds with a single session. Session number is indicated at the left of each row, with italicized font indicating a high-external-return session.

C Screenshots

Introduction

Identification Number: 1

Welcome and thank you for participating. Just for agreeing to participate you will automatically be given $10.00 as a "thank you" payment. Anything else you earn today will be in addition to this.

Your name will never be recorded in this study. Instead, you will be known by your Identification Number that appears above.

In this experiment you will be asked to complete seven experimental tasks. In each task you will be assigned to the same group, consisting of all 8 participants in this room. Your earnings in each task will depend on your own decision, as well as on the decisions of all 7 other participants. After all tasks are completed, one of the tasks will be randomly chosen. You will only be paid for the chosen task. It makes good sense, therefore, to complete each task as though it will actually be chosen. Further instructions will be provided at the beginning of each task.

Before we proceed, we will take a few minutes to introduce you to your group members. We will first ask Number 1 to stand and say to everyone "Hello. I am Number 1." We'll then ask Number 2 to do similarly, and will repeat this for everyone.

Begin now with Number 1.

Please wait until all introductions are done.

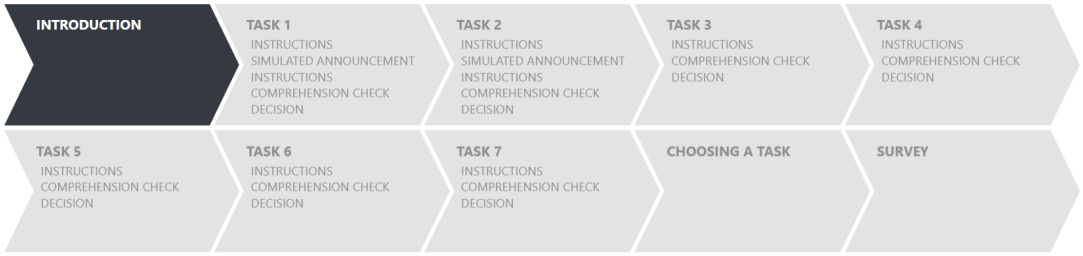

Task 1

Identification Number: 1

Instructions

You have been given $10.00 to divide between a personal account and a group account. Note that only whole-dollar divisions are allowed. Every dollar you allocate to your personal account will earn you one dollar. However, every dollar allocated to the group account (either by you or by any other group member) will earn $0.30 for the subject who allocated it, and $0.50 for each of the other group members.

Therefore, your earnings from this task will be:

The number of dollars you allocate to your personal account

+

0.3 times the number of dollars you allocate to the group account

+

0.5 times the number of dollars all 7 other group members allocate to the group account.

Note that regardless of what the other group members choose to do, the more you allocate to your personal account, the greater will be your earnings from this task. However, the group as a whole will benefit from every dollar allocated to the group account.

Example: Suppose each group member allocated $0.00 to their personal account and $10.00 to the group account. Then each group member would earn ($0.00 + (0.3 * $10.00) + (0.5 * 7 * $10.00) =) $38.00.

Example: Suppose each group member allocated $10.00 to their personal account and $0.00 to the group account. Then each group member would only earn ($10.00 + (0.3 * $0.00) + (0.5 * 7 * $0.00) =) $10.00.

Example: Suppose each of the other group members allocated all $10.00 to the group account, while you allocated all $10.00 to your personal account. Then you would earn ($10.00 + (0.3 * $0.00) + (0.5 * 7 * $10.00) =) $45.00, while each of the other group members would earn (0 + (0.3 * $10.00) + (0.5 * 6 * $10.00) =) $33.00.

Example: Suppose each of the other group members allocated all $10.00 to their personal account, while you allocated all $10.00 to the group account. Then you would earn (0 + (0.3 * $10.00) + (0.5 * 7 * 0) =) $3.00, while each of the other group members would earn ($10.00 + (0.3 * 0) + (0.5 * $10.00) =) $15.00.

Example: Suppose you allocated $8.00 to your personal account and $2.00 to the group account, and all 7 other group members allocated a total of $47.00 to the group account. Then you would earn ($8.00 + (0.3 * $2.00) + (0.5 * $47.00) =) $32.10.

Example: Suppose you allocated $1.00 to your personal account and $9.00 to the group account, and all 7 other group members allocated a total of $36.00 to the group account. Then you would earn ($1.00 + (0.3 * $9.00) + (0.5 * $36.00) =) $21.70.

Next

Task 1

Identification Number: 1

Instructions

Announcements

You will make your allocation decision in private, and will receive no feedback until the very end of the experiment. After you have completed all of the tasks, one of the tasks will be randomly chosen. If Task 1 is chosen, then an announcement will be made about each group member's Selected Allocation in this task. This announcement may or may not be the same as the group member's Selected Allocation, and will be determined as follows:

Each of you will be asked to spin a virtual roulette wheel like this:

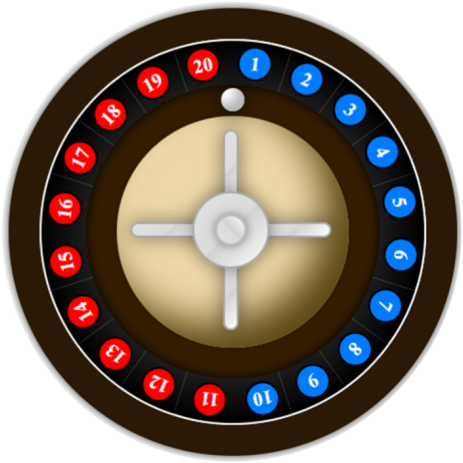

- If your spin result is one of the following:
  1 2 3 4 5 6 7 8 9 10
  then your **Selected Allocation** will be announced.
- However, if your spin result is one of the following:
  11 12 13 14 15 16 17 18 19 20
  then a **random allocation** will be announced instead of your Selected Allocation. This random allocation will be determined by asking you to roll a virtual 11-sided die numbered 0-10. Nobody but you will see that a die is being rolled, or its result. The result of this die roll will be your Announced Allocation to the group account. For example, if the result of the die roll is 5, then your Announced Allocation to the group account will be $5.00.

Therefore, if this task is chosen at the end of the experiment:

- *Everyone in this room will know the Announced Allocation of each group member.*
- *No one will be told whether this Announced Allocation is the Selected Allocation decision or a random one.*
- *Payments will be assigned according to each member's Selected Allocation, not the Announced Allocation.*
- *Everyone in this room will know his/her payment only after leaving the experiment.*

Announcements will be made at the end of the experiment by displaying on everyone's screen something similar to this:

Chosen task: 1

Odds of announcing Selected Allocation: 10 in 20 (50%)

Odds of announcing random allocation: 10 in 20 (50%)

Subject	**Personal account**	**Group account**
Subject 1	$2.00	$8.00
Subject 2	$9.00	$1.00
Subject 3	$5.00	$5.00
Subject 4	$6.00	$4.00
Subject 5	...and so forth.	

We will also read each subject's Announced Allocation out loud, and ask you each to stand up and face the other subjects while an announcement is made about your Announced Allocation.

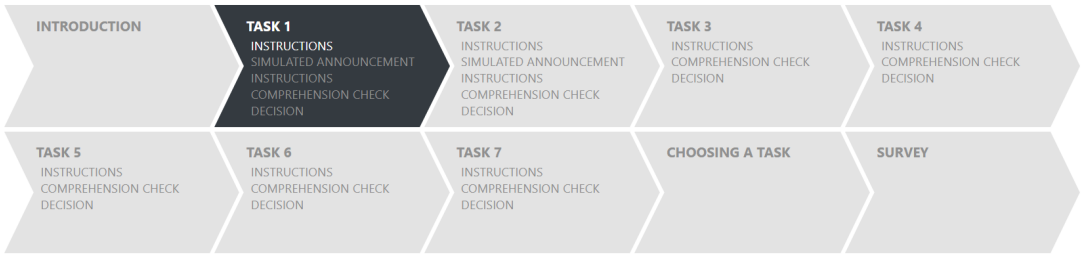

Next

Task 1

Identification Number: 1

Simulated Announcement

To make sure everyone understands how the announcements will be determined if Task 1 is chosen, we will now run a simulation. In the simulation, instead of allowing each person to choose his or her Selected Allocation, a random division of the $10.00 will be selected for each of you. We will ask you to imagine that this is your Selected Allocation in the simulated announcement.

Simulated announcements will then be determined as explained before. That is:

- You will be asked to spin a virtual roulette wheel.
- The spin result will determine whether your Selected Allocation will be announced or whether a random allocation will be announced, in which case you will be asked to roll a virtual 11-sided die that will determine your Announced Allocation to the group account.
- The Announced Allocations will then be displayed on everyone's screen.
- We will read each subject's Announced Allocation aloud while this subject stands up and faces the other subjects.

After the simulation, you will be asked to answer a series of comprehension questions, and then you will be asked to make your decision for this task.

Keep in mind that the allocations in this simulation are imaginary and will thus have no effect on your actual earnings from this task.

Next

Task 1

Identification Number: 1

Simulated Announcement

Imagine that in this simulated task your Selected Allocation was to allocate $4.00 to your personal account and $6.00 to the group account. Imagine further that we are now at the end of the experiment and this task is chosen.

To determine whether your simulated Selected Allocation will be announced or whether a random allocation will be announced in this simulated announcement, please spin the following roulette wheel by clicking the button "Click To Spin".

Remember that if your spin result is one of the following:

1 2 3 4 5 6 7 8 9 10

then your **Selected Allocation** will be announced.

However, if your spin result is one of the following:

11 12 13 14 15 16 17 18 19 20

then a **random allocation** will be announced instead of your Selected Allocation.

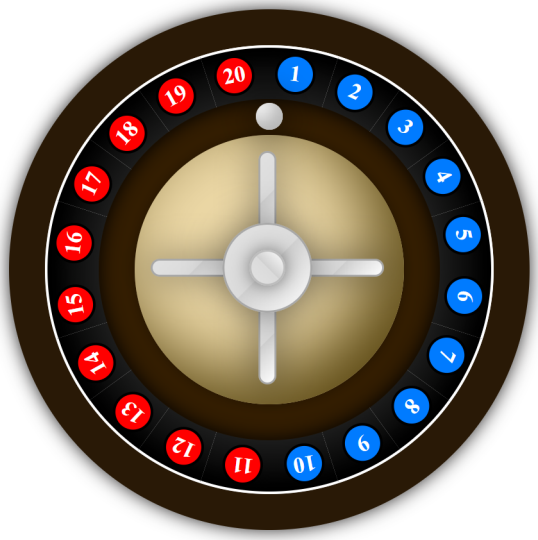

Click To Spin

Keep in mind that the allocations in this simulation are imaginary and will thus have no effect on your actual earnings from this task.

Task 1

Identification Number: 1

Simulated Announcement

To determine the random allocation that will be announced, please roll the following 11-sided die by clicking the button "Click To Roll".

Remember that the result of this die roll will be your Announced Allocation to the group account. For example, if the result of the die roll is 5, then your Announced Allocation to the group account will be $5.00.

Click To Roll

Keep in mind that the allocations in this simulation are imaginary and will thus have no effect on your actual earnings from this task.

Task 1

Identification Number: 1

Simulated Announcement

Your simulated Selected Allocation was to allocate $4.00 to your personal account and $6.00 to the group account.

Your simulated Announced Allocation will be $2.00 to your personal account and $8.00 to the group account.

Keep in mind that the allocations in this simulation are imaginary and will thus have no effect on your actual earnings from this task.

Next

Task 1

Identification Number: 1

Simulated Announcement

Chosen task: 1

Odds of announcing Selected Allocation: 10 in 20 (50%)

Odds of announcing random allocation: 10 in 20 (50%)

Subject	Personal account	Group account
Subject 1	$2.00	$8.00
Subject 2	$6.00	$4.00
Subject 3	$10.00	$0.00
Subject 4	$8.00	$2.00
Subject 5	$0.00	$10.00
Subject 6	$5.00	$5.00
Subject 7	$7.00	$3.00
Subject 8	$2.00	$8.00

Keep in mind that the allocations in this simulation are imaginary and will thus have no effect on your actual earnings from this task.

Identification Number: 1

End of Simulation; Start of Actual Task.

Next

Task 1

Identification Number: 1

Recall the instructions and announcement procedure for Task 1.

Announcements

You will make your allocation decision in private, and will receive no feedback until the very end of the experiment. After you have completed all of the tasks, one of the tasks will be randomly chosen. If Task 1 is chosen, then an announcement will be made about each group member's Selected Allocation in this task. This announcement may or may not be the same as the group member's Selected Allocation, and will be determined as follows:

Each of you will be asked to spin a virtual roulette wheel like this:

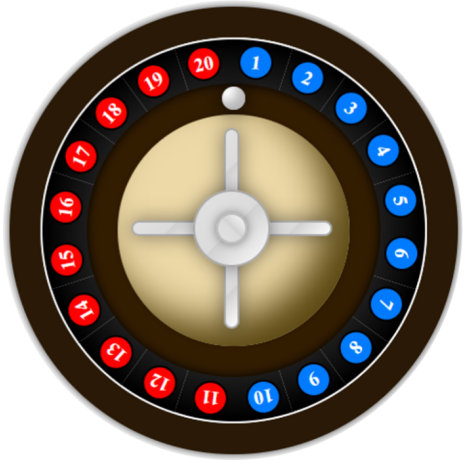

- If your spin result is one of the following:
  1 2 3 4 5 6 7 8 9 10
  then your **Selected Allocation** will be announced.
- However, if your spin result is one of the following:
  11 12 13 14 15 16 17 18 19 20
  then a **random allocation** will be announced instead of your Selected Allocation. This random allocation will be determined by asking you to roll a virtual 11-sided die numbered 0-10. Nobody but you will see that a die is being rolled, or its result. The result of this die roll will be your Announced Allocation to the group account. For example, if the result of the die roll is 5, then your Announced Allocation to the group account will be $5.00.

Therefore, if this task is chosen at the end of the experiment:

- *Everyone in this room will know the Announced Allocation of each group member.*
- *No one will be told whether this Announced Allocation is the Selected Allocation decision or a random one.*
- *Payments will be assigned according to each member's Selected Allocation, not the Announced Allocation.*
- *Everyone in this room will know his/her payment only after leaving the experiment.*

Announcements will be made at the end of the experiment by displaying on everyone's screen something similar to this:

Chosen task: 1

Odds of announcing Selected Allocation: 10 in 20 (50%)

Odds of announcing random allocation: 10 in 20 (50%)

Subject	Personal account	Group account
Subject 1	$2.00	$8.00
Subject 2	$9.00	$1.00
Subject 3	$5.00	$5.00
Subject 4	$6.00	$4.00
Subject 5	...and so forth.	

We will also read each subject's Announced Allocation out loud, and ask you each to stand up and face the other subjects while an announcement is made about your Announced Allocation.

Next

Task 1

Identification Number: 1

Comprehension Check

Suppose you allocated $3.00 to your personal account and $7.00 to the group account, and all other group members allocated a total of $24.00 to the group account. How much would you earn if this task is chosen at the end of the experiment (in addition to the $10.00 show-up payment)?

$

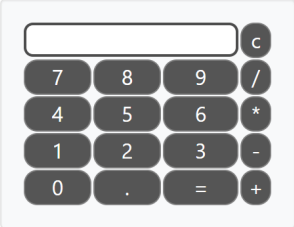

Click here to see the instructions again

Next

Task 1

Identification Number: 1

Comprehension Check

$6.00 is incorrect.

The question was: Suppose you allocated $3.00 to your personal account and $7.00 to the group account, and all other group members allocated a total of $24.00 to the group account. How much would you earn if this task is chosen at the end of the experiment (in addition to the $10.00 show-up payment)?

The correct answer is: $17.10.

Explanation: If this task is chosen at the end of the experiment, then your earnings would be:

The number of dollars you allocate to your personal account
+
0.3 times the number of dollars you allocate to the group account
+
0.5 times the number of dollars all 7 other group members allocate to the group account.

Therefore, you would earn: ($3.00 + 0.3 * $7.00 + 0.5 * $24.00 =) $17.10.

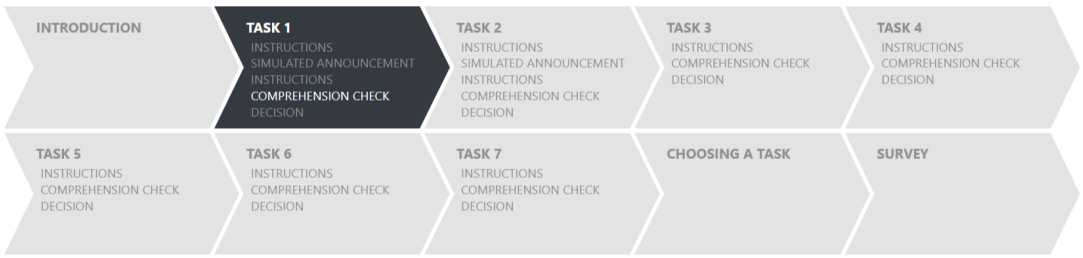

Next

Task 1

Identification Number: 1

Comprehension Check

Suppose you allocated $6.00 to your personal account and $4.00 to the group account, and all other group members allocated a total of $43.00 to the group account. How much would you earn if this task is chosen at the end of the experiment (in addition to the $10.00 show-up payment)?

$

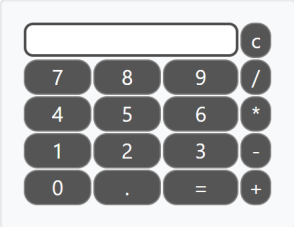

Click here to see the instructions again

Next

Task 1

Identification Number: 1

Comprehension Check

Correct!

The question was: Suppose you allocated $6.00 to your personal account and $4.00 to the group account, and all other group members allocated a total of $43.00 to the group account. How much would you earn if this task is chosen at the end of the experiment (in addition to the $10.00 show-up payment)?

The correct answer is: $28.70.

Explanation: If this task is chosen at the end of the experiment, then your earnings would be:

> The number of dollars you allocate to your personal account
>
> +
>
> 0.3 times the number of dollars you allocate to the group account
>
> +
>
> 0.5 times the number of dollars all 7 other group members allocate to the group account.

Therefore, you would earn: ($6.00 + 0.3 * $4.00 + 0.5 * $43.00 =) $28.70.

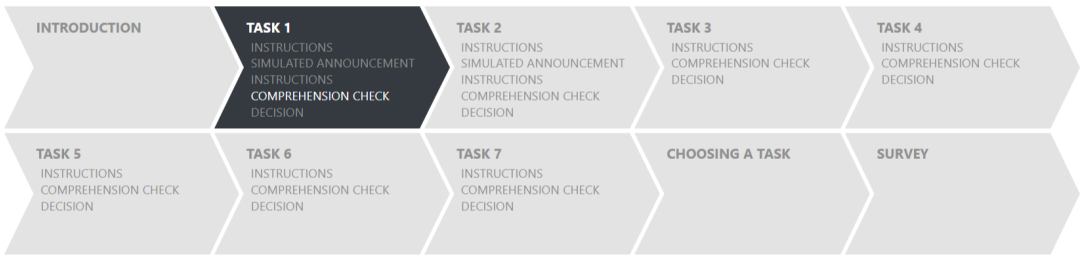

Next

Task 1

Identification Number: 1

Comprehension Check

Imagine that you allocated $5.00 to the group account, that the current task is chosen at the end of the experiment, and that your spin result is 8.

(a) What would be your Announced Allocation to the group account?

- $8.00
- $2.00
- It could be any whole-dollar amount from $0.00 to $10.00, depending on the result of an 11-sided die roll
- $5.00

(b) How much of your endowment would actually be allocated to the group account?

- $5.00
- It could be any whole-dollar amount from $0.00 to $10.00, depending on the result of an 11-sided die roll
- $8.00
- $2.00

Click here to see the instructions again

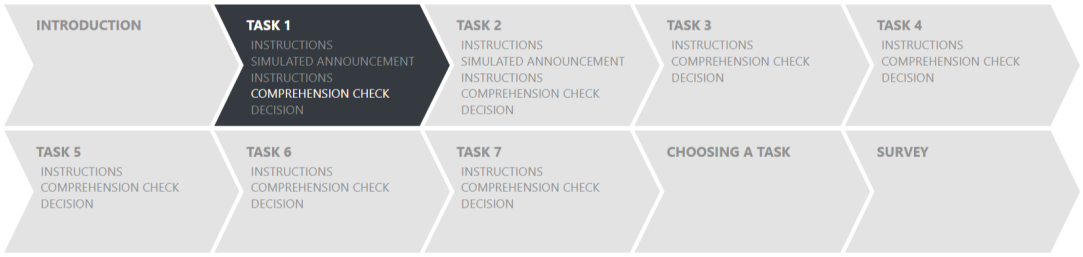

Next

Task 1

Identification Number: 1

Comprehension Check

Correct answers!

The question was:

Imagine that you allocated $5.00 to the group account, that the current task is chosen at the end of the experiment, and that your spin result is 8.

(a) What would be your Announced Allocation to the group account?
(b) How much of your endowment would actually be allocated to the group account?

The correct answers are:

(a) $5.00.
(b) $5.00.

Explanation:

If this this task is chosen at the end of the experiment, then your Announced Allocation will be determined as follows:

You will be asked to spin a virtual roulette wheel, such as the one you have seen in the simulation round.

- If your spin result is one of the following:
  1 2 3 4 5 6 7 8 9 10
  then your **Selected Allocation** will be announced.
- However, if your spin result is one of the following:
  11 12 13 14 15 16 17 18 19 20
  then a **random allocation** will be announced instead of you your Selected Allocation. This random allocation will be determined by asking you to roll a virtual 11-sided die numbered 0-10. The result of this die roll will be your Announced Allocation to the group account.

If your spin result were 8, then your Selected Allocation would be announced. Therefore, your Announced Allocation to the group account would be $5.00 if this were the amount you had chosen to allocate to the group account.

In addition, your Selected Allocation will determine how much of your endowment would actually be allocated to the group account. Therefore, if you chose to allocate $5.00 of your endowment to the group account, then this is the amount that truly would be allocated to the group account.

Next

Task 1

Identification Number: 1

Comprehension Check

Imagine that you allocated $5.00 to the group account, that the current task is chosen at the end of the experiment, that your spin result is 12, and that the result of your die roll is 2.

(a) What would be your Announced Allocation to the group account?

- It could be any whole-dollar amount from $0.00 to $10.00, depending on the result of an 11-sided die roll
- $2.00
- $8.00
- $5.00

(b) How much of your endowment would actually be allocated to the group account?

- $8.00
- It could be any whole-dollar amount from $0.00 to $10.00, depending on the result of an 11-sided die roll
- $5.00
- $2.00

Click here to see the instructions again

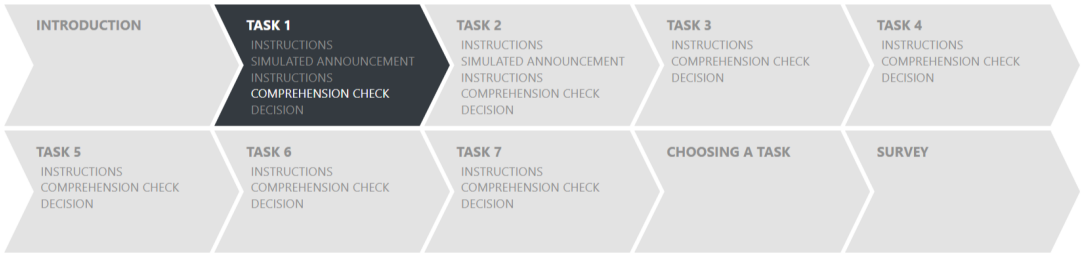

Next

Task 1

Identification Number: 1

Comprehension Check

Correct answers!

The question was:

Imagine that you allocated $5.00 to the group account, that the current task is chosen at the end of the experiment, that your spin result is 12, and that the result of your die roll is 2.

(a) What would be your Announced Allocation to the group account?
(b) How much of your endowment would actually be allocated to the group account?

The correct answers are:

(a) $2.00.
(b) $5.00.

Explanation:

If this this task is chosen at the end of the experiment, then your Announced Allocation will be determined as follows:

You will be asked to spin a virtual roulette wheel, such as the one you have seen in the simulation round.

- If your spin result is one of the following:
  1 2 3 4 5 6 7 8 9 10
  then your **Selected Allocation** will be announced.
- However, if your spin result is one of the following:
  11 12 13 14 15 16 17 18 19 20
  then a **random allocation** will be announced instead of you your Selected Allocation. This random allocation will be determined by asking you to roll a virtual 11-sided die numbered 0-10. The result of this die roll will be your Announced Allocation to the group account.

If your spin result were 12, a random allocation would be announced instead of your Selected Allocation. If the result of your die roll were 2, your Announced Allocation to the group account would be $2.00.

However, your Selected Allocation will determine how much of your endowment would actually be allocated to the group account. Therefore, if you chose to allocate $5.00 of your endowment to the group account, then this is the amount that truly would be allocated to the group account.

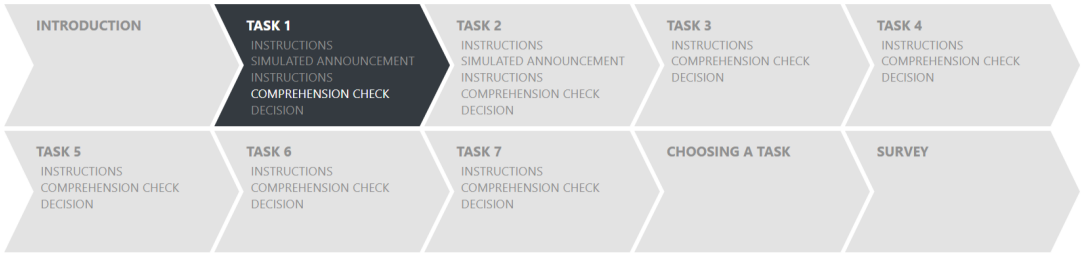

Next

Task 1

Identification Number: 1

Comprehension Check

Imagine that the current task is chosen at the end of the experiment. What would be the probability that we announce a group member's random allocation?

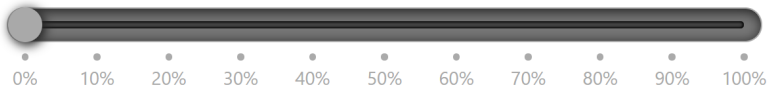

0% chance of random allocation

Click here to see the instructions again

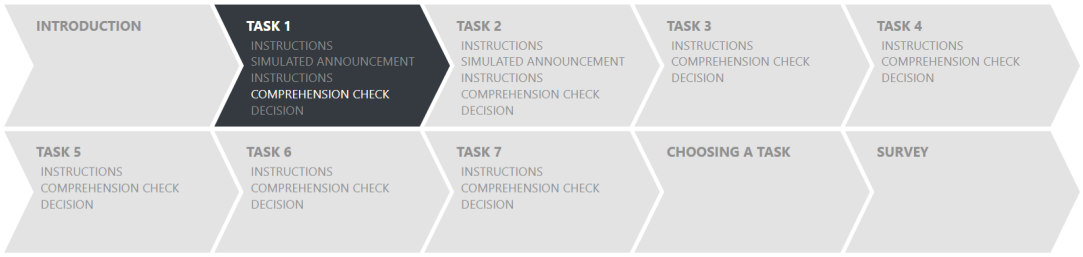

Next

Task 1

Identification Number: 1

Comprehension Check

25% is incorrect.

The question was: Imagine that the current task is chosen at the end of the experiment. What would be the probability that we announce a group member's random allocation?

The correct answer is: 50%.

Next

Task 1

Identification Number: 1

Comprehension Check

Imagine that the current task is chosen at the end of the experiment. How would everyone's payment be determined (in addition to the $10.00 show-up payment)?

- Based on each group member's Announced Allocation
- Everyone would receive $10.00
- Based on each group member's Selected Allocation
- At random

Click here to see the instructions again

Next

Task 1

Identification Number: 1

Comprehension Check

Correct!

The question was: Imagine that the current task is chosen at the end of the experiment. How would everyone's payment be determined (in addition to the $10.00 show-up payment)?

The correct answer is: Based on each group member's Selected Allocation.

Next

Task 1

Identification Number: 1

Your Decision

You will now make your decision for Task 1. If this task is the one that is chosen at random at the end of the experiment, then these decisions will be used to determine your payment.

Given a 50% chance that your Selected Allocation will be announced, and a 50% chance that a random allocation will be announced, how would you like to divide your $10.00 between the two accounts?

Dollars in personal account:

$ - Each dollar earns $1.00 for you and $0.00 for each of the other group members

Dollars in group account:

$ - Each dollar earns $0.30 for you and $0.50 for each of the other group members

Note that the amounts you enter must be whole numbers that sum to $10.00.

Next

Identification Number: 1

You have now completed Task 1. We now move on to Task 2.

Next

Task 2

Identification Number: 1

Instructions

Announcements

This task will be much the same as the previous task, except that the odds of announcing a group member's Selected Allocation (if this task is selected at the end of the experiment), will be the following:

Odds of announcing Selected Allocation: 1 in 20 (5%)
Odds of announcing random allocation: 19 in 20 (95%)

The announcements will be determined as follows:

Each of you will be asked to spin a virtual roulette wheel, such as the one you have seen.

- If your spin result is:
  1
  then your **Selected Allocation** will be announced.
- However, if your spin result is one of the following:
  2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
  then a **random allocation** will be announced instead of your Selected Allocation. This random allocation will be determined by asking you to roll a virtual 11-sided die numbered 0-10. Nobody but you will see that a die is being rolled, or its result. The result of this die roll will be your Announced Allocation to the group account. For example, if the result of the die roll is 5, then your Announced Allocation to the group account will be $5.00.

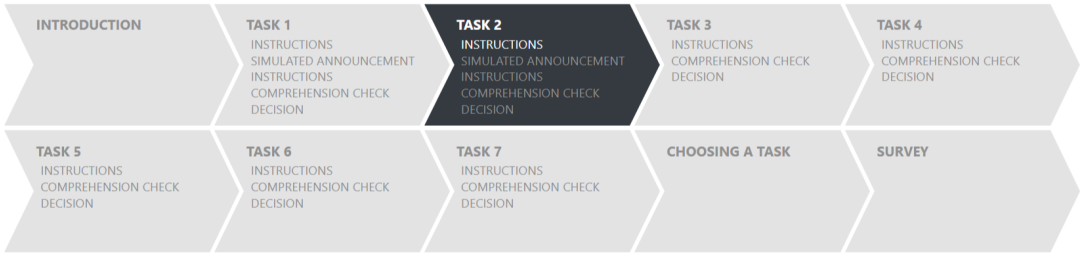

Next

⋮

And so forth.

Choosing a Task

Identification Number: 1

We will now choose one of the 7 tasks at random, and use only that task to determine all participants' payments and announcements. After the announcements have been made, we will ask you to complete a survey. After everyone has finished answering this survey, we will hand out the payments, and you will be free to leave.

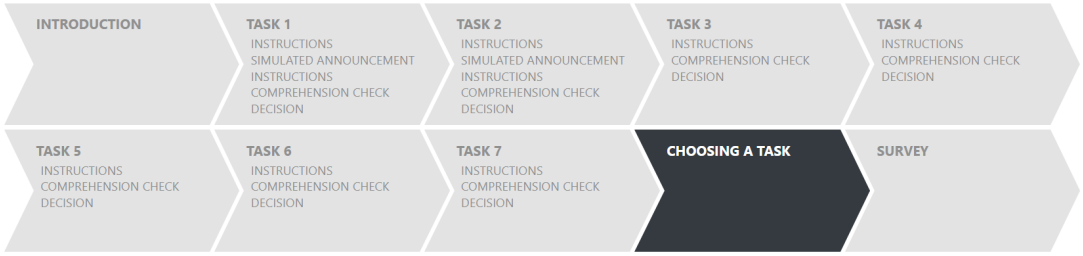

Next

Your Announced Allocation

Identification Number: 1

The chosen task is 2. Therefore, the following holds:

Odds of announcing Selected Allocation: 1 in 20 (5%)
Odds of announcing random allocation: 19 in 20 (95%)

In this task you chose to allocate $5.00 to your personal account and $5.00 to the group account.

To determine whether your Selected Allocation will be announced or whether a random allocation will be announced, please spin the following roulette wheel by clicking the button "Click To Spin".

If your spin result is:

1

then your **Selected Allocation** will be announced.

However, if your spin result is one of the following:

2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20

then a **random allocation** will be announced instead of you your **Selected Allocation**.

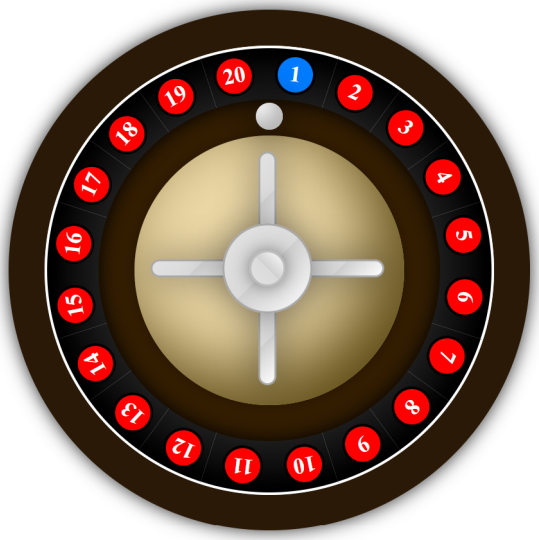

Click To Spin

Your Announced Allocation

Identification Number: 1

To determine the random allocation that will be announced, please roll the following 11-sided die by clicking the button "Click To Roll".

The result of this die roll will be your announced allocation to the group account. For example, if the result of the die roll is 5, then your announced allocation to the group account will be $5.00.

Click To Roll

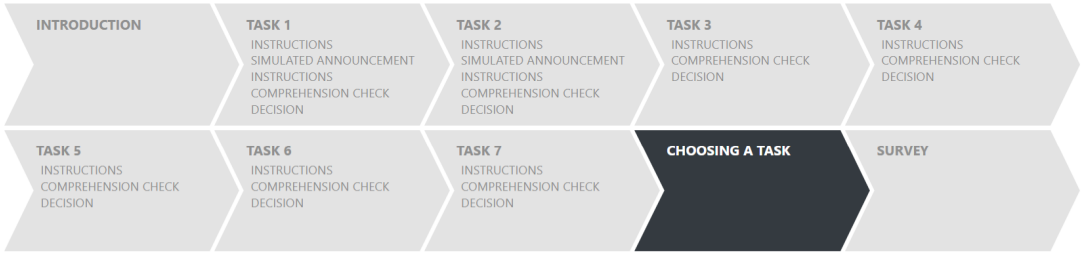

Your Announced Allocation

Identification Number: 1

Your Selected Allocation in Task 2 was to allocate $5.00 to your personal account and $5.00 to the group account.

Your Announced Allocation will be $4.00 to your personal account and $6.00 to the group account.

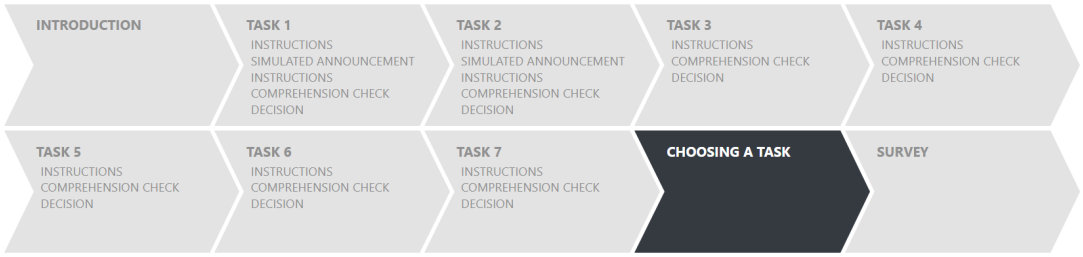

Next

Announcements

Identification Number: 1

Chosen Task: 2

Odds of announcing Selected Allocation: 1 in 20 (5%)

Odds of announcing random allocation: 19 in 20 (95%)

Subject	Personal account	Group account
Subject 1	$4.00	$6.00
Subject 2	$3.00	$7.00
Subject 3	$9.00	$1.00
Subject 4	$2.00	$8.00
Subject 5	$6.00	$4.00
Subject 6	$3.00	$7.00
Subject 7	$3.00	$7.00
Subject 8	$7.00	$3.00

Identification Number: 1

We are almost done with the experiment. On the next few screens, for statistical-analysis purposes, we are going to ask you questions about yourself. Please answer each question as best you can.

Next

Survey

Identification Number: 1

Here are a number of characteristics that may or may not apply to you. For example, do you agree that you are someone who likes to spend time with others? Please write a number next to each statement to indicate the extent to which you agree or disagree with that statement.

I See Myself as Someone Who..

	Disagree strongly	Disagree a little	Neither agree nor disagree	Agree a little	Agree strongly
1. Is talkative	○	○	○	○	○
2. Tends to find fault with others	○	○	○	○	○
3. Does a thorough job	○	○	○	○	○
4. Is depressed, blue	○	○	○	○	○
5. Is original, comes up with new ideas	○	○	○	○	○
6. Is reserved	○	○	○	○	○
7. Is helpful and unselfish with others	○	○	○	○	○
8. Can be somewhat careless	○	○	○	○	○
9. Is relaxed, handles stress well	○	○	○	○	○
10. Is curious about many different things	○	○	○	○	○
11. Is full of energy	○	○	○	○	○
12. Starts quarrels with others	○	○	○	○	○
13. Is a reliable worker	○	○	○	○	○
14. Can be tense	○	○	○	○	○
15. Is ingenious, a deep thinker	○	○	○	○	○
16. Generates a lot of enthusiasm	○	○	○	○	○
17. Has a forgiving nature	○	○	○	○	○
18. Tends to be disorganized	○	○	○	○	○
19. Worries a lot	○	○	○	○	○
20. Has an active imagination	○	○	○	○	○
21. Tends to be quiet	○	○	○	○	○
22. Is generally trusting	○	○	○	○	○
23. Tends to be lazy	○	○	○	○	○
24. Is emotionally stable, not easily upset	○	○	○	○	○

25. Is inventive	○	○	○	○	○
26. Has an assertive personality	○	○	○	○	○
27. Can be cold and aloof	○	○	○	○	○
28. Perseveres until the task is finished	○	○	○	○	○
29. Can be moody	○	○	○	○	○
30. Values artistic, aesthetic experiences	○	○	○	○	○
31. Is sometimes shy, inhibited	○	○	○	○	○
32. Is considerate and kind to almost everyone	○	○	○	○	○
33. Does things efficiently	○	○	○	○	○
34. Remains calm in tense situations	○	○	○	○	○
35. Prefers work that is routine	○	○	○	○	○
36. Is outgoing, sociable	○	○	○	○	○
37. Is sometimes rude to others	○	○	○	○	○
38. Makes plans and follows through with them	○	○	○	○	○
39. Gets nervous easily	○	○	○	○	○
40. Likes to reflect, play with ideas	○	○	○	○	○
41. Has few artistic interests	○	○	○	○	○
42. Likes to cooperate with others	○	○	○	○	○
43. Is easily distracted	○	○	○	○	○
44. Is sophisticated in art, music, or literature	○	○	○	○	○

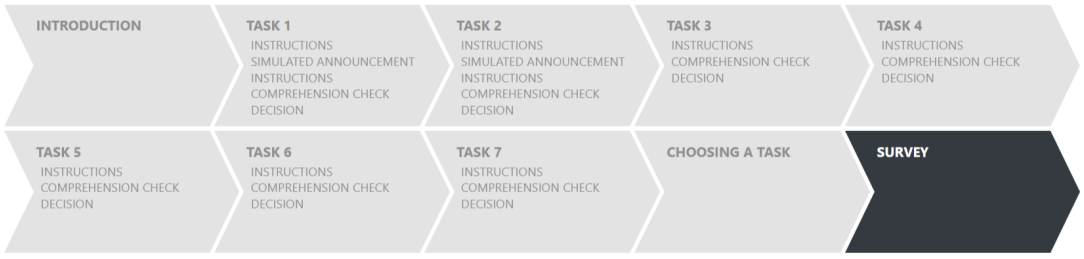

Next

Survey

Identification Number: 1

Read each of the following statements carefully and indicate how characteristic it is of you.

	Not at all characteristic of me	Slightly characteristic of me	Moderately characteristic of me	Very characteristic of me	Extremely characteristic of me
1. I worry about what other people will think of me even when I know it doesn't make any difference.	○	○	○	○	○
2. I am unconcerned even if I know people are forming an unfavorable impression of me.	○	○	○	○	○
3. I am frequently afraid of other people noticing my shortcomings.	○	○	○	○	○
4. I rarely worry about what kind of impression I am making on someone.	○	○	○	○	○
5. I am afraid that others will not approve of me.	○	○	○	○	○
6. I am afraid that people will find fault with me.	○	○	○	○	○
7. Other people's opinions of me do not bother me.	○	○	○	○	○
8. When I am talking to someone, I worry about what they may be thinking about me.	○	○	○	○	○
9. I am usually worried about what kind of impression I make.	○	○	○	○	○
10. If I know someone is judging me, it has little effect on me.	○	○	○	○	○
11. Sometimes I think I am too concerned with what other people think of me.	○	○	○	○	○
12. I often worry that I will say or do the wrong things.	○	○	○	○	○

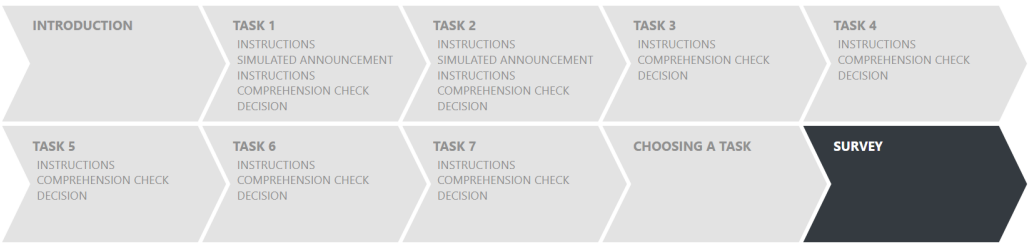

Next

Survey

Identification Number: 1

Please state the extent to which each of the following statements is true of you.

	Not at all true of me	1	2	3	4	5	6	7	Very true of me
1. When I see people I do not know feeling sad, I feel a need to reach out to them.		○	○	○	○	○	○	○	
2. I spend a lot of time concerned about the well-being of humankind.		○	○	○	○	○	○	○	
3. When I hear about someone (a stranger) going through a difficult time, I feel a great deal of compassion for him or her.		○	○	○	○	○	○	○	
4. It is easy for me to feel the pain (and joy) experienced by others, even though I do not know them.		○	○	○	○	○	○	○	
5. If I encounter a stranger who needs help, I would do almost anything I could to help him or her.		○	○	○	○	○	○	○	
6. I feel considerable compassionate love for people from everywhere.		○	○	○	○	○	○	○	
7. I would rather suffer myself than see someone else (a stranger) suffer.		○	○	○	○	○	○	○	
8. If given the opportunity, I am willing to sacrifice in order to let the people from other places who are less fortunate achieve their goals.		○	○	○	○	○	○	○	
9. I tend to feel compassion for people even though I do not know them.		○	○	○	○	○	○	○	
10. One of the activities that provides me with the most meaning to my life is helping others in the world who need help.		○	○	○	○	○	○	○	
11. I would rather engage in actions that help others, even though they are strangers, than engage in actions that would help me.		○	○	○	○	○	○	○	

12. I often have tender feelings toward people (strangers) when they seem to be in need.	○ ○ ○ ○ ○ ○ ○
13. I feel a selfless caring for most of mankind.	○ ○ ○ ○ ○ ○ ○
14. I accept others whom I do not know even when they do things I think are wrong.	○ ○ ○ ○ ○ ○ ○
15. If a person (a stranger) is troubled, I usually feel extreme tenderness and caring.	○ ○ ○ ○ ○ ○ ○
16. I try to understand rather than judge people who are strangers to me.	○ ○ ○ ○ ○ ○ ○
17. I try to put myself in a stranger's shoes when he or she is in trouble.	○ ○ ○ ○ ○ ○ ○
18. I feel happy when I see that others (strangers) are happy.	○ ○ ○ ○ ○ ○ ○
19. Those whom I encounter through my work and public life can assume that I will be there if they need me.	○ ○ ○ ○ ○ ○ ○
20. I want to spend time with people I don't know well so that I can help enrich their lives.	○ ○ ○ ○ ○ ○ ○
21. I very much wish to be kind and good to fellow human beings.	○ ○ ○ ○ ○ ○ ○

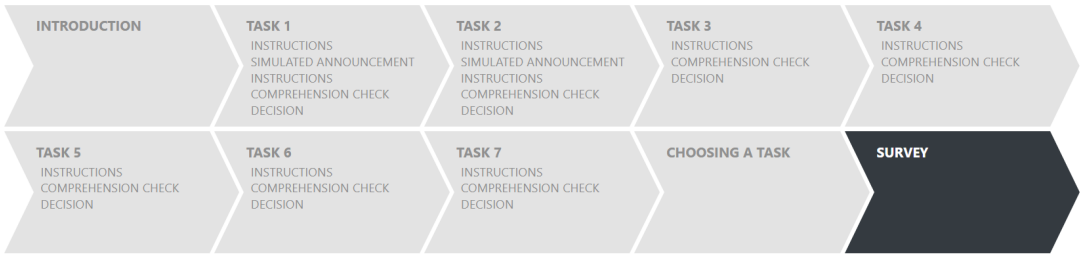

Next

Survey

Identification Number: 1

Please state the extent to which you agree or disagree with each of the following statements.

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
Privacy laws should be strengthened to protect personal privacy.	○	○	○	○	○
People need legal protection against misuse of personal data.	○	○	○	○	○
If I were to write a constitution today, I would probably add privacy as a fundamental right.	○	○	○	○	○
When I share the details of my personal life with somebody, I often worry that he/she will tell those details to other people.	○	○	○	○	○
I am concerned that people around me know too much about me.	○	○	○	○	○
I am concerned with the consequences of sharing identity information	○	○	○	○	○
I worry about sharing information with more people than I intend to.	○	○	○	○	○
If somebody is not careful about protecting their own privacy, I cannot trust them about respecting mine.	○	○	○	○	○
If I am to enjoy some privacy in my life, I need my friends to be careful about protecting their privacy as well.	○	○	○	○	○
I could never trust someone as my confidant if they go around sharing details about their own private lives.	○	○	○	○	○
The level of privacy that I can enjoy depends on the extent to which people around me protect their own privacy.	○	○	○	○	○
It is important for me to respect the privacy of individuals, even if they are not careful about protecting their own privacy.	○	○	○	○	○
I value other people's privacy as much as I value mine.	○	○	○	○	○
Even when somebody is not careful about his/her privacy, I do my best to respect that person's privacy.	○	○	○	○	○
I always do my best not to intrude into other people's private lives.	○	○	○	○	○
Respect for others' privacy should be an important priority in social relations.	○	○	○	○	○

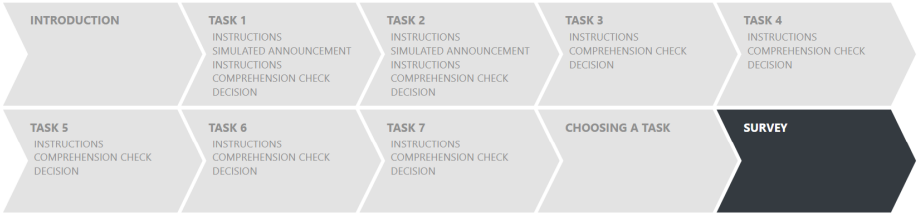

Next

Survey

Identification Number: 1

What is your gender?

- Female
- Male
- Other

Are you of Hispanic origin or descent?

- Yes
- No

What race do you consider yourself?

- White
- Black/African American
- Asian or Pacific Islander
- American Indian/Native American
- Other race (please specify)

In what year were you born?

What is the highest level of education you have completed?

- Middle school or less
- Some high school
- High school diploma
- GED (HS Equivalent)
- Some college, but did not finish
- Two-year college degree/Associate degree/A.A./A.S.
- Four-year college degree/B.A./B.S.
- Some graduate school
- Master's degree (MA/MS/MBA/MFA/MDiv)
- Advanced degree (PhD/MD/JD)

What is your major?

Thinking about economic issues, which of the following best describes your attitudes?

- Very liberal
- Liberal
- Slightly liberal
- Moderate
- Slightly conservative
- Conservative
- Very conservative

Thinking about social issues, which of the following best describes your attitudes?

- Very liberal
- Liberal
- Slightly liberal
- Moderate
- Slightly conservative
- Conservative
- Very conservative

Do you consider yourself a...

- Republican
- Democrat
- Independent
- Moderate
- Other
- None of the above

Do you have any comments about the experiment? Please provide them in the text box below. We would love to hear any feedback, suggestions, and thoughts you may have.

We are particularly curious to know: How did you decide how to allocate the money?

What do you think the experiment is about?

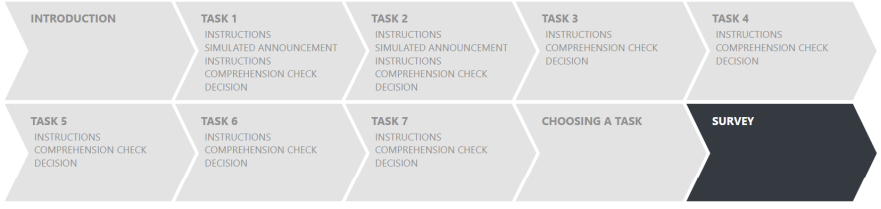

Next

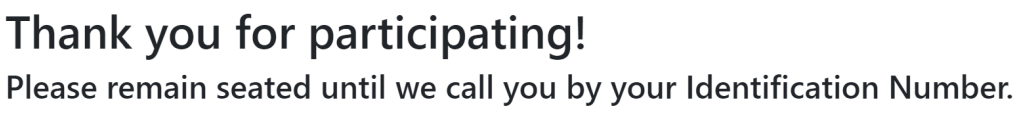

References

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.

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.

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.

Eckel, Catherine C., and Philip J. Grossman. 2006. "Subsidizing charitable giving with rebates or matching: Further laboratory evidence." *Southern Economic Journal*, 72(4): 794–807.

Eckel, Catherine C., and Philip J. Grossman. 2008. "Subsidizing charitable contributions: a natural field experiment comparing matching and rebate subsidies." *Experimental Economics*, 11(3): 234–252.

Eckel, Catherine C., and Philip J. Grossman. 2017. "Comparing rebate and matching subsidies controlling for donors' awareness: Evidence from the field." *Journal of Behavioral and Experimental Economics*, 66: 88–95.

Gandullia, Luca. 2019. "The price elasticity of warm-glow giving." *Economics Letters*, 182: 30–32.

Gandullia, Luca, and Emanuela Lezzi. 2018. "The price elasticity of charitable giving: New experimental evidence." *Economics Letters*, 173: 88–91.

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.

Huck, Steffen, and Imran Rasul. 2011. "Matched fundraising: Evidence from a natural field experiment." *Journal of Public Economics*, 95(5-6): 351–362.

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.

Karlan, Dean, and John A. List. 2007. "Does price matter in charitable giving? Evidence from a large-scale natural field experiment." *American Economic Review*, 97(5): 1774–1793.

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

Meer, Jonathan. 2014. "Effects of the price of charitable giving: Evidence from an online crowdfunding platform." *Journal of Economic Behavior & Organization*, 103: 113–124.

Peloza, John, and Piers Steel. 2005. "The price elasticities of charitable contributions: a meta-analysis." *Journal of Public Policy & Marketing*, 24(2): 260–272.

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.

Scharf, Kimberley, and Sarah Smith. 2015. "The price elasticity of charitable giving: does the form of tax relief matter?" *International Tax and Public Finance*, 22(2): 330–352.

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