Revealed Privacy Preferences:
Are Privacy Choices Rational?

Yi-Shan Lee[†] and Roberto A. Weber[†]

Abstract

We investigate the extent to which tradeoffs involving the sharing of personal information exhibit consistency with an underlying rational preference for privacy. In an experiment, people engage in tradeoffs across two domains of personal information, allowing us to classify whether their choices satisfy the Generalized Axiom of Revealed Preference. Sixty-three percent of subjects act consistently with a rational preference ordering when allocating privacy levels, despite substantial heterogeneity of privacy attitudes. Individuals who are inconsistent when engaging in such privacy tradeoffs exhibit substantially more costly preference reversals when pricing the sharing of their personal information. Our results imply that allocating privacy property rights by monetizing personal information can have distinct monetary welfare consequences for people with different degrees of rationality in their underlying ability to make sensible tradeoffs involving personal information sharing. We also provide evidence that preferences elicited over choices in our experiment correlate with real-world privacy behaviors.

Keywords: information sharing, rationality, experiment, GARP, personal data

[†] Lee: CUHK Business School, Chinese University of Hong Kong, Shatin, Hong Kong.
Weber: Department of Economics, University of Zurich, Bluemlisalpstrasse 10, Zurich, Switzerland.
E-mail: yishanlee@cuhk.edu.hk (Lee) and roberto.weber@econ.uzh.ch (Weber).
We are thankful for comments and suggestions from participants at several seminars and conferences: Foundations of Utility and Risk Workshop, Society for the Advancement of Economic Theory Conference, Stanford Institute of Theoretical Economics (SITE), Security and Human Behavior Workshop, Annual Digital Information Policy Scholar Conference, Symposium: Privacy at Michigan, ESA World Meeting, seminar participants at University of Amsterdam, University of Vienna, George Mason University, Chinese University of Hong Kong, Law and Economics Seminar, Empirical Legal Research Workshop, Evidence-based Economics Summer Meeting and Experimental and Behavioral Economics Seminar. We are especially grateful to Alessandro Acquisti for comments on an early version of our manuscript. Yi-Shan Lee gratefully acknowledges financial support from the Swiss National Science Foundation (Grant 174630) and the Forschungskredit of the University of Zurich (Grant No. FK-17-017).

1 Introduction

Privacy policies play a critical role in the information economy, influencing the effectiveness of online marketing (Goldfarb & Tucker, 2011), crowdfunding (Burtch, Ghose & Wattal, 2015), market concentration in the technology sector (Peukert et al., 2020), and technology adoption in health (Miller & Tucker, 2009; Miller & Tucker, 2018). How to design efficient policies for maintaining privacy, while not harming the potential benefits from the voluntary exchange of information, is a source of considerable debate. While the European Union considers personal data protection a fundamental human right and establishes strict regulations (European Parliament, Council of the European Union, 2016), the United States takes a more self-regulatory approach that emphasizes free exchange to allocate property rights over personal data (The Privacy Office, US Department of Homeland Security, 2009). The need for and characteristics of policies that protect individual privacy are critically related to such individuals' abilities to appropriately manage the sharing of their personal information. If individuals cannot behave rationally when sharing their personal data, such violations of rationality can justify the need for privacy regulations.[1] However, given the potential heterogeneity of privacy attitudes and the complexity of tradeoffs involved in personal data sharing decisions, it can be challenging to make a definitive judgment about whether individuals are capable of sharing or concealing personal information rationally.

Basic notions of economic rationality provide a means for objectively measuring whether individuals' privacy choices are consistent with models based on utility maximization. Such measures identify the extent to which choices satisfy the necessary properties to be represented by a complete and transitive preference ordering. Applying this standard of rationality to the tradeoffs that individuals make in simple privacy-related contexts can yield insights into whether they can manage privacy in consistent and sensible manners in much more complex real-world settings.

In this study, we investigate the degree to which privacy choices are consistent with a rational preference ordering by applying a "revealed privacy preference" approach to two kinds of privacy decisions: (i) sharing different kinds of personal information and (ii) tradeoffs between privacy control and money. Both kinds of decisions are valuable for understanding

[1] Theoretical research on privacy in economics typically assumes rational decision making regarding the use of personal information, focusing instead, for example, on informational externalities produced by individuals' sharing decisions (Choi, Jeon & Kin, 2019; Acemoglu et al., 2021) and implications of consumer privacy choices for online markets (Casadesus-Masanell & Hervas-Drane, 2015).

whether people can manage the sharing of their personal information. The first kind of choice sheds light on how people decide privacy levels for different forms of personal information and whether they appear to do so as if they are maximizing a coherent set of preferences that can be characterized by a utility function. The second kind of choice quantifies the cost of decision errors when people must consider trading off privacy for money.[2] Thus, our study provides information regarding the feasibility of individuals efficiently allocating their personal information by giving them property rights over their information. The measurement of simple notions of rationality in these two kinds of privacy decisions provides insights into how to design privacy policies, by yielding information on how frequently, and at which point, inconsistencies in decisions involving privacy may arise.

Furthermore, we also study the connection between individuals' privacy choices in the above two stylized laboratory contexts and their privacy behavior outside the laboratory. We do so by eliciting incentivized measures of how much personal information people share in several real-world domains and investigating the extent to which these behavioral concerns for privacy reflect the preferences we elicit using our laboratory measures. Thus, we can relate both the level of privacy that individuals select in our laboratory choices and the rationality of such choices to their information sharing outside the laboratory.

An important property of our study is that the use of subjects' data and the consequences of personal information sharing are clear and understandable. There is no uncertainty about how the data will be used or what will happen due to its use. Hence, this setup presents a useful starting point for understanding the rationality of privacy preferences in a setting where the tradeoffs are simple and explicitly specified and all consequences are known. If subjects fail to manage personal data sensibly in this setting, it is unlikely they will be able to do so when the perceived sharing consequences are ambiguous and more complicated in the real world. In this sense, the level of rationality we observe in this study can be thought of as an upper bound on the extent to which people can rationally manage their privacy.

[2] Many studies study the value of personal information and information use in online markets (e.g., Lewis & Anderson, 2005; Korunovska & Bauer, 2012; Carrascal et al., 2013; Danezis, Simonite, 2014; Spiekermann, Kummer & Schulte, 2019). Our study adds to this literature by quantifying errors due to inconsistency when individuals engage in tradeoffs involving their personal information. Our study is thus relevant for better understanding the welfare implications of (models of) privacy policies, including utilizing differential privacy mechanisms to determine the optimal level of of privacy protection (Abowd & Schmutte, 2019), taxes and mediated data sharing to ameliorate the negative externality of oversharing personal data (Acemoglu et al., 2021) and allocating privacy property rights to consumers to maximize social welfare given the nonrivalry of data (Jones & Tonetti, 2020).

In the first part of this study, we begin by testing whether individuals control and share different personal data items in a systematic manner consistent with a stable underlying preference. At the core of rational privacy behavior lies the deliberate concealment or revelation of personal information (Posner, 1981), and we test this using a simple choice setting. We do this by asking individuals to choose their most preferred privacy level, measured by the number of people who will view two kinds of "reports" containing different forms of personal information: the subject's body fat composition and a measure of the subject's intelligence. We present these choices in a sequence of decision problems that vary in the relative "prices" of revealing the two kinds of information. The variation in the budget sets and the corresponding relative prices across decisions in this part provide a stringent test of whether choices satisfy the properties of utility maximization.

More precisely, in Part 1 of our study we classify people according to general privacy attitudes over these types of information and evaluate the rationality of individual privacy choices by their consistency with the General Axiom of Revealed Preference (GARP). We find that most (68%) of our subjects value the privacy of both information items as goods, and more than 72 percent of subjects of this private type make choices that are entirely consistent with the maximization of an underlying convex preference. The privacy attitudes of other people are public (10%), meaning they value sharing both kinds of information, or item-dependent (22%). Despite the heterogeneity of privacy attitudes, we find more than 63 percent of individuals in our pooled sample share and conceal their personal data in a perfectly consistent manner. Hence, about two-thirds of our subjects' privacy preferences across different information domains can be represented as a well-behaved preference ordering consistent with utility maximization. This proportion of subjects acting consistently with well-behaved privacy preferences is comparable to the proportion of subjects acting consistently with well-behaved risk preferences with the analogous budget sets, elicited in a follow-up study. In addition, this proportion is comparable to those in several previous experimental studies testing consistency with GARP when making other types of tradeoffs.

In the second part of this study, we utilize the revealed privacy preferences from Part 1 to test the extent to which people leave money on the table when privacy decisions involve tradeoffs between privacy control and money. This provides insights into the efficiency and implications of tackling privacy issues by employing market mechanisms involving tradable privacy property rights. We compare the willingness-to-accept (WTA) for the revelation of bundles of information, including some preferable bundles, selected by subjects from the budget sets in the first part of the experiment, and others unchosen from the same sets. A preference reversal in this context implies that a subject requires a lower price for selling their information according to a bundle, *X*, than for an alternative bundle that their choices in Part 1 reveal as directly preferred to *X*. Interestingly, 47 percent of the subjects we classified as rational in Part 1 exhibit such reversals, while 65 percent of those classified as irrational do so. Thus, a substantial proportion of subjects stand to suffer utility loss in privacy tradeoffs that involve money, even in our very simple choice setting.

We also quantify the degree of utility loss from mistakes people make when mapping their revealed choice preferences into monetary valuations, which we refer to as "monetary-value noise" (MVN). We find substantial heterogeneity across subjects in this measure. Most interestingly, people whose privacy preference in Part 1 violate GARP exhibit significantly larger average MVN—260 percent larger than the MVN of individuals we classify as rational in Part 1. This result implies that allocating privacy property rights by monetizing or licensing personal information can have distinct monetary welfare consequences for people with different degrees of rationality in their underlying ability to make sensible tradeoffs involving personal information sharing.

However, we also show that despite heterogeneity in MVN and WTA across individuals, at the group level, private types have a significantly higher mean willingness-to-accept (WTA) for sharing their data than public types. Thus, monetary valuations like WTA seem to tell us something meaningful about privacy preferences. Nevertheless, the heterogeneity and noise in WTA values in our study imply that such valuations should be used with caution.

The third part of this study involves an exploratory investigation into the relationship between our primary privacy measures in the lab and other measures of privacy preferences and behavior outside the laboratory. Since the scale of collection and sharing of personal data has increased significantly, it is essential to understand the extent to which this creates economic value that individuals can manage in a rational and beneficial manner. Thus, we explore whether individuals we classify differently based on their privacy orientations and rationality in Part 1 of our experiment exhibit different real-world behavioral patterns when managing their personal information. We investigate this link between the lab and the field by eliciting information on subjects' daily-life sharing of personal details using incentivized questions. Specifically, we elicit concrete measures of how much information subjects share online and with firms and provide them with incentives for truthfully revealing instances of information sharing. We then test whether individuals with different privacy attitudes and degrees of rationality behave differently when sharing their personal data.

While these exploratory results must be interpreted cautiously, we find evidence of a relationship between our laboratory privacy measures and some individual privacy behavior in the field. Specifically, factor analysis identifies two primary domains of daily-life information sharing in our data: social-networking use (SN) and the exchange of personal information for money or services (EMS). We find that SN correlates strongly with privacy attitudes but not with rationality, while EMS has a modest correlation with rationality. We cautiously interpret these correlational results as suggesting that individuals' privacy preferences have some underlying stability, but also that the relationship between such preferences and behavior may be domain-specific.

Finally, we also use our laboratory measures of privacy preferences to test the validity of a widely used self-reported measure of privacy attitudes. We find that private types who prefer to show their personal information to fewer viewers in the experiment report greater concern for privacy, in general. Thus, we provide some validation for this question as a measure of privacy preferences.

Broadly, our study constitutes a significant step forward in measuring privacy preferences and understanding the extent to which individuals engage in rational tradeoffs when making decisions involving the use of their personal information. Employing GARP as a necessary condition for rationality, we find significant heterogeneity in how rational people can be when managing their privacy, even in the very simple context we study. We also find that people who fail to satisfy this necessary condition for rationality suffer substantial monetary losses when trading off their privacy control for money. Finally, simple laboratory choices measuring privacy preferences have some predictive value for real-world privacy behaviors and general privacy attitudes. Our results suggest that careful consideration should be given to the heterogeneity of rationality and its monetary consequences when designing privacy policies and that simple stylized choices can be used to understand important aspects of how people approach the management and sharing of their personal information.

The remainder of the article is structured as follows. Section 2 reviews related literature. Section 3 describes the experimental design. Section 4 reports results of the three parts of our study: 4.1 studies the rationality of privacy choices when deciding between bundles of information sharing, 4.2 presents the consistency between these privacy choices and their monetary equivalents, and 4.3 further analyzes the consistency between laboratory privacy measures and privacy behaviors in the field and general privacy attitudes. Section 5 concludes.

2 Related Literature

This project is closely related to debates regarding the need for privacy protection policies, at the core of which are often arguments regarding the rationality of individuals' decisions involving tradeoffs between privacy protection and the benefits of sharing personal information. Pioneered by Chicago School scholars George Stigler (Stigler, 1961, 1962, 1980) and Richard Posner (Posner, 1978, 1981), the anti-regulation perspective highlights that privacy protection is redistributive and creates market inefficiency. However, the rationality assumptions underlying the Chicago privacy models and the welfare implications were concurrently questioned by Hirshleifer (1971, 1980), and later by many other scholars. For a comprehensive summary of this literature, see Acquisti et al. (2016).

Surprisingly, although there is mixed theoretical and empirical evidence on whether privacy protections are beneficial or harmful for individuals and society across different situations, most studies make welfare implications in a context where individuals' abilities to manage personal data sharing rationally are left largely undiscussed, and the existence of a rational privacy preference is untested. Among the few exceptions, the most related studies to this project are Taylor (2004) and Villas-Boas (2004). Both papers find the ability to control and share personal data rationally or even strategically is decisive for the welfare effects of privacy protection regulations. Inspired by the importance of rationality in the literature on the privacy protection debate, this project aims to fill important gaps in the literature by being the first to directly test individuals' rationality when making two essential privacy tradeoffs: between the sharing of different kinds of personal information and between privacy control and monetary payoffs.

The literature on privacy choices involving different information items has shown that people have different preferred privacy levels for different items, and that the willingness to allow different numbers of viewers can serve as a measure of privacy concerns. Schudy and Utikal (2015) validate that privacy, defined according to the number of strangers who receive a piece of personal information, is valued as an economic good. In particular, individuals' willingness to share their personal data weakly decreases in the number of recipients and varies across different information items. This work provides a foundation for our first tradeoff task, in which we ask whether people trade off privacy levels across different personal information domains in a systematic way consistent with an underlying preference assigning (dis)utility to the sharing of personal information. Empirically, social networking site users already report that if they could grant access to limited viewers, they would have shared approximately half the unshared content (Sleeper et al., 2013). This finding suggests that people may have different sensitivity to privacy levels for different information items. Our first tradeoff task aims to contribute to this line of literature by further testing if the preferred privacy levels across different information items differ systematically across different decision problems in response to implicit "prices," and thus rationally in the economic sense.[3]

On the other hand, the literature on tradeoffs between privacy control and monetary benefits shows that the monetary value of privacy control is not only heterogeneous across individuals but also malleable. In Schudy and Utikal (2015), the monetary value of privacy across individuals exhibits a high degree of heterogeneity: around 12 percent of participants always share personal information to earn money, whereas about 20 percent of participants never share, and the decisions of the rest of participants depend on monetary payoffs. However, the monetary value of privacy control not only varies individually, but also varies depending on the elicitation method (Benndorf and Normann, 2017), the endowment and the order of privacy choices (Acquisti et al., 2013), and the salience of privacy control (Tsai et al., 2011; Beresford et al., 2012). Surprisingly, even as monetary value varies, the share of individuals sharing a particular item remains relatively stable (Benndorf and Normann, 2017). This phenomenon raises the possibility that people, in principle, know whether they want to share a particular information item, but are unsure how to map this decision into the monetary domain. Although many studies have investigated how people value various types of privacy control, in different scenarios using various methodologies (for a detailed survey, see Kokolakis (2017)), taking the monetary value as a good measure for privacy preferences requires an auxiliary assumption—specifically, that people can sensibly translate their privacy preferences into monetary values—that has yet to be thoroughly tested. The validation of this assumption is essential for the legitimacy of utilizing market mechanisms to allocate privacy, as proposed in a substantial literature on privacy protection (Laudon, 1996; Samuelson, 2000; Schwartz, 2004; Varian, 2009; Ghosh and Roth, 2015).

[3] This rationality test can also support the utility-theoretic interpretation of a key privacy notion for privacy protection technology that emerged from computer science, differential privacy (see Dwork (2008) for an excellent overview).

This study also contributes to a better understanding of the relationship between rational economic and behavioral approaches to privacy choices. There are diverse streams of empirical studies on the influence of behavioral and psychological phenomena on privacy choices; see Acquisti et al. (2015, 2020) for a good overview. Although law articles often quote behavioral research to contend that the current practice of privacy self-management "does not provide people meaningful control over their data" (Solove, 2013, p.1880), economists may note that specific behavioral phenomena, in the context of privacy choices, do not mean privacy preferences are not rationalizable. For example, the documented effect of default choices on privacy settings (Stutzman et al., 2012; Acquisti and Ralph, 2006) and the effects of the endowed privacy level on the monetary value of privacy control (Acquisti et al., 2013) may both be rationalized by context-dependent preferences with a reasonable degree of flexibility in defining the domain of preferences (Kalai et al., 2002). Compared to the existing empirical and experimental studies, the dataset produced by this project—including the elicitation of tradeoffs across varying domains and behaviors—aims to disentangle privacy preferences systematically and to test if individuals can make privacy tradeoffs rationally within and across different domains when the choice environment is simple, when people are well informed and when the tradeoffs in which they engage are salient. Thus, while we go beyond existing tests of the rationality of privacy preferences, our analysis is one in which the expectation should be of a high level of consistent and rational behavior, if there is much hope of finding it in richer and more complex settings.

3 The Experiment

3.1 Design Overview

The central part of our experiment consists of an assessment process to generate the personal information and three stages in which subjects make decisions. During the assessment process, each subject generates two different forms of personal information: a body-composition score and an intelligence score. Subjects are told that this information may subsequently be used to generate two hard copy "assessment reports," each including a subject's name, signature, picture, and one of the two above pieces of personal information.

Stages 1 and 2 involve making decisions that may lead to subjects forfeiting their property rights over the personal information in these two reports. The choices in these stages involve "report viewing combinations" of their personal information. In each combination, a given number of strangers will view their body-composition score, and another given number of strangers will view their intelligence score. As a result of their choices in Stages 1 and 2, a particular bundle may be implemented and the corresponding reports shown to the specified number of individuals. A detailed timeline of the experiment is shown in Table 1.

Table 1: Experimental procedures

Stages		Steps	
0	Instructions and assessment process	1	Read written instructions and provide consent
		2	Take a profile picture
		3	Measure body fat composition
		4	Measure intelligence
		5	Receive two assessment reports and two copies of the profile picture in envelopes
1	Tradeoffs between different information items	1	Detailed instruction for stage 1
		2	4 privacy choices for privacy-attitude classification (Stage 1A)
		3	16 privacy-tradeoff choices between combinations of numbers of viewers for assessment reports (Stage 1B)
2	Tradeoffs between money and privacy	1	Detailed instruction for stage 2
		2	Indicate WTA for showing reports to strangers: 40 combinations (20 chosen choices + 16 perturbed choices + 4 pre-determined combinations)
3	Real-life privacy behavior	1	Detailed instruction for stage 3
		2	Measure 1: public personal information items on Facebook
		3	Measure 2: use of club cards and store memberships
		4	Measure 3: active social-media accounts
		5	Measure 4: use of location-sharing apps
		6	Measure 5: number of Facebook friends
		7	Measure 6: number of Facebook friends living in Zurich
4	Survey	1	Personal characteristics and open questions

In Stage 1, subjects are confronted with 20 budget sets, each presenting several possible bundles from which a subject can choose. A subject selects 20 possible bundles, thereby indicating their preferred use of their personal information given the feasible budgets. At the end of Stage 1, the computer selects another 20 possible uses. Then in Stage 2, subjects decide, for each of the above 40 possible uses, for what range of prices they are willing to sell their personal information so that it is used in the described manner. Choices in these two stages determine how the personal information in the two assessment reports is used and also the main portion of the final payment. Finally, in Stage 3, subjects answer an incentivized questionnaire about their daily-life personal information sharing behavior. Details of the procedures for each stage are described in the next section.

The experiment lasted 2 hours, and subjects received a participation fee of CHF 40 (Swiss francs, with CHF 1 ≈ USD 1). However, they could add to this payment during the experiment. The final payments consisted of anonymous bonus payments for intelligence tests[4] that ranged from CHF 0 to CHF 6 and main payments, including the participation fee, that ranged from CHF 40 to CHF 95.5.

3.2 Stage 0: Assessment Process

At the beginning of the experiment, subjects provided three pieces of personal information: a profile picture taken by the experimenter, a measure of body fat composition and an intelligence score.[5] Subjects completed this assessment privately, and their scores were only observed by one experimenter, who performed no other function in the session than administering the assessment and recording the scores. This experimenter had no access to subjects' names or any other personally-identifying information.

All information was recorded in hard copies by the experimenter. Then, the subject received hard copies of "assessment reports" containing these data in closed envelopes. The reports contained the following information: a subject's portrait photograph and either the subject's body fat or the subject's intelligence. At this point, the subject's name was not on the report, though there was a blank space where this information could be recorded. Meanwhile, subjects were aware that they could make choices in later stages of the experiment for their assessment reports to be shown to "viewers"—randomly recruited strangers from the University of Zurich and the Swiss Federal Institute of Technology (ETH). These viewers would see a series of assessment reports in exchange for a small payment. Subjects were also informed that, if their assessment reports would be shown to viewers, viewers would see the reports on a computer display individually, but the reproduction or transmission of the personal information would not be possible.

3.3 Stage 1: Specifying uses of personal information

In this stage, each subject chose 20 possible uses for their assessment reports from feasible sets, by indicating their most preferred "report-viewing combination" using sliders in each decision problem. A combination (x, y) refers to some number, x, of people who will view the

[4] Subjects received this anonymous bonus payment in private, without filling a receipt, and from an experimenter who had no access to their names. This maintained privacy over their intelligence scores.

[5] The measure of body fat was obtained by having subjects hold an electronic sensor in the presence of an experimenter; this sensor reported a score that was viewed by the subject and confirmed by the experimenter. The intelligence score was obtained by subjects answering a 12-item computerized version of Raven's progressive matrices for which they could obtain more money for each correct answer.

body composition report and another number, y, of people who will view the intelligence report. Subjects were aware that the computer would then choose another 20 combinations and that these 40 combinations would subsequently be the only possible uses of their reports after Stage 1.

In the first four decision problems of Stage 1 (Stage 1A), subjects created a combination of viewers by choosing the number of viewers for each assessment report independently, without making tradeoffs between the privacy levels of the two assessment reports. Two of these questions were single-dimensional in the sense that subjects chose freely from the range [0,125] for one report, and the number of viewers of the other report was fixed at zero. These two questions were used as practice rounds to familiarize subjects with the interface. The choice set of the other two questions was two-dimensional, as shown by the two shaded squares in Figure 1(a). One question has a range [0,80] for both reports, and the other has a range of [45,125] for both reports. On these two questions, subjects could choose any point in the square choice set, in increments of 5. Thus, these choices allow us to observe what direction a subject prefers in a two-dimensional space of possible viewing combinations. We later use the decisions in these two-dimensional questions to classify subjects into different privacy attitudes: private (a preference for fewer viewers on both dimensions), public (more viewers on both dimensions) or item-dependent (more viewers on one dimension and fewer on the other).

In the remaining 16 decision problems of Stage 1 (Stage 1B), subjects made tradeoffs between privacy levels for the two assessment reports. In each problem, subjects chose a combination consisting of some number of viewers for each report, requiring tradeoffs between the viewers for the two reports at “prices” that were fixed for the decision problem. The full set of budgeted privacy levels is shown in Figure 1(b). These budget sets allow us to identify GARP violations for each possible kind of privacy attitude.

Figure 2(a) shows an example of the interface used to represent the budget sets to subjects. In this example, a subject may explore all linear combinations on a budget line connecting two extreme combinations: combination A (100 viewers of the body composition report and 0 viewers of the intelligence report) and combination B (0 viewers of the body composition report and 50 viewers of the intelligence report.) by using a slider that controls both slider bars simultaneously. The implicit price for this budget line equals 2 body composition viewers for 1 intelligence viewer. The subject then chose his or her favorite

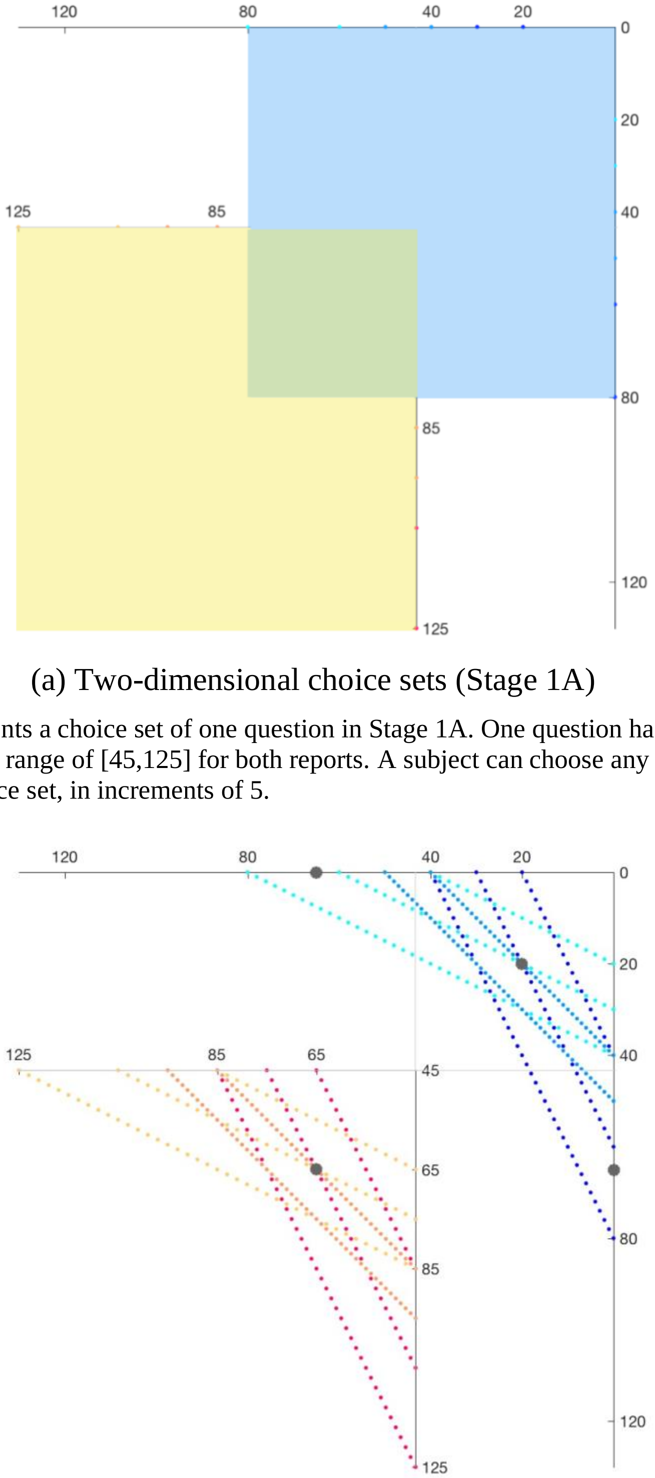

(a) Two-dimensional choice sets (Stage 1A)

Notes: Each square represents a choice set of one question in Stage 1A. One question has a range [0,80] for both reports, and the other has a range of [45,125] for both reports. A subject can choose any combination of viewers for both reports in the choice set, in increments of 5.

(b) Downward-sloping budget sets (Stage 1B) and preset bundles (Stage 2)

Notes: Each downward-slopping line represents a choice set of one question in Stage 1B; with one dot represents an available combination of viewers for both reports. The large four grey dots represent the preset bundles for eliciting willingness-to-pay in Stage 2.

Figure 1: Choice sets and preset combinations used in Stages 1 and 2

bundled numbers of viewers (e.g., 36 viewers of the body composition report and 32 viewers of the intelligence report) on the budget line.[7]

At the end of Stage 1, subjects had selected 20 combinations in total. This included 2 from one-dimensional problems, 2 from the independent two-dimensional privacy-attitude measures and 16 from downward-sloping budget sets that allow us to identify GARP violations.

3.4 Stage 2: Reviewing and selling the use of personal information

At the beginning of Stage 2, the computer generated 16 additional combinations—one from each of the downward-sloping budget sets—for each of our subjects. Specifically, for each budget set in Figure 1(b), the computer randomly chose one of the two points that was exactly five units away from the combination chosen by a subject.[9] Thus, each of these 16 combinations represents a bundle that was not selected when a different bundle was chosen. This procedure we use for selecting these combinations keeps the difference in the number of viewers between perturbed choices and chosen choices constant across subjects.

In Stage 2, subjects then traded off privacy control against monetary payoffs for 40 report-viewing combinations, including all 20 report-viewing combinations they had selected in Stage 1, the 16 additional combinations selected by the computer as described above and 4 additional preset combinations that were the same for all participants: (20,20), (65,65), (0,65) and 65,0), as shown by the four grey dots in Figure 1(b). These additional 4 combinations allow us to make direct comparisons across subjects' valuations for a fixed set of bundles.

Subjects first learned the fixed range of possible prices (from 0 CHF to 40 CHF) from which the computer would offer them a randomly selected price for selling their personal information in each decision problem. For each possible use of their information (a combination of viewers for each report), they then decided the lowest acceptable price for which they were willing to sell their personal information and allow it to be viewed by people according to that report-viewing combination. Subjects always had the option to state a price of "more than 40 CHF," to indicate that they were unwilling to sell the personal information at any of the possible prices. Figure 2(b) shows an example of the interface of this Stage.

[7] We can use the test developed by Bronars (1987) to evaluate how stringent our revealed preference test is for each type. This Bronars power test provides the probability that a subject who chose purely randomly would have a revealed GARP violation given the possible violations defined by a set of feasible choices. Using simulations with 10,000 randomly choosing individuals, we find the set of choices in Figure 1(b) has a Bronars power of 0.9996 for the private type, of 0.9419 for the public type, and of 0.9999 for the item-dependent types (when requiring no GARP violations in both directions). The design thus has sufficient power to make very unlikely the possibility that we mistakenly classify individuals as rational.

[9] If the subject chose an extreme combination, the computer used the only point that was five units away.

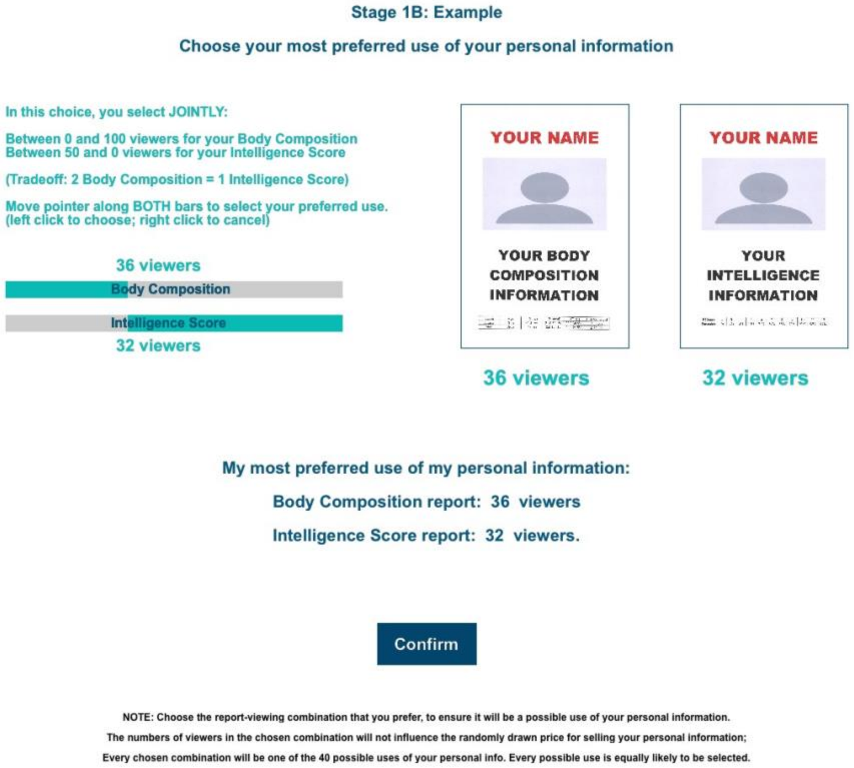

Stage 1B: Example

Choose your most preferred use of your personal information

In this choice, you select JOINTLY:

Between 0 and 100 viewers for your Body Composition
Between 50 and 0 viewers for your Intelligence Score

(Tradeoff: 2 Body Composition = 1 Intelligence Score)

Move pointer along BOTH bars to select your preferred use.
(left click to choose; right click to cancel)

36 viewers

32 viewers

My most preferred use of my personal information:

Body Composition report: 36 viewers

Intelligence Score report: 32 viewers.

NOTE: Choose the report-viewing combination that you prefer, to ensure it will be a possible use of your personal information.
The numbers of viewers in the chosen combination will not influence the randomly drawn price for selling your personal information;
Every chosen combination will be one of the 40 possible uses of your personal info. Every possible use is equally likely to be selected.

(a) Privacy allocations (Stage 1)

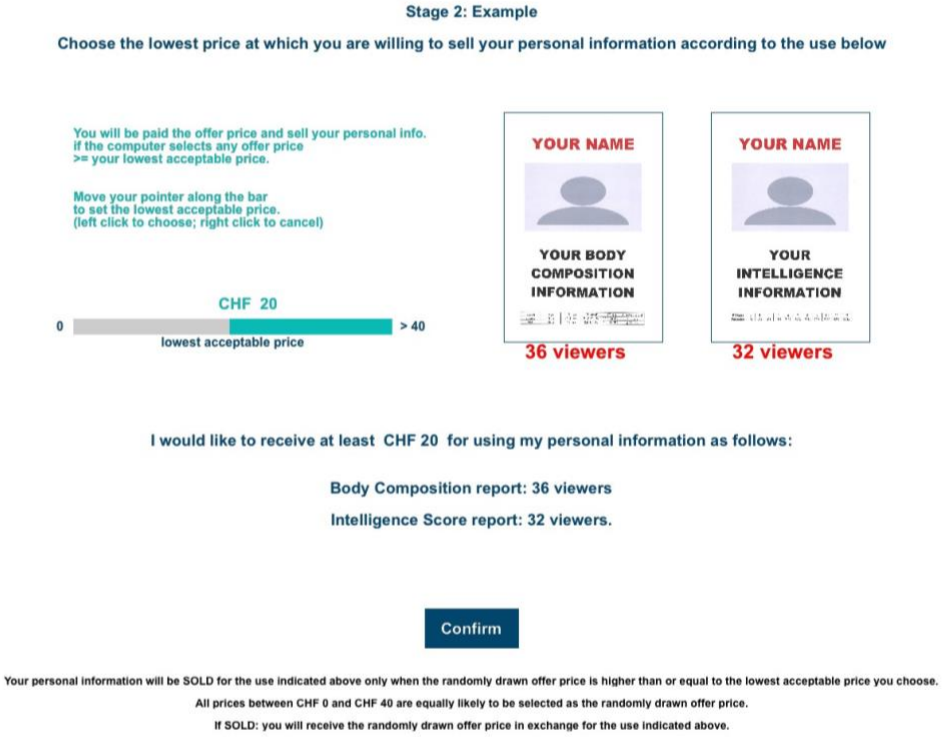

Stage 2: Example

Choose the lowest price at which you are willing to sell your personal information according to the use below

You will be paid the offer price and sell your personal info.
if the computer selects any offer price
>= your lowest acceptable price.

Move your pointer along the bar
to set the lowest acceptable price.
(left click to choose; right click to cancel)

CHF 20

I would like to receive at least CHF 20 for using my personal information as follows:

Body Composition report: 36 viewers

Intelligence Score report: 32 viewers.

Your personal information will be SOLD for the use indicated above only when the randomly drawn offer price is higher than or equal to the lowest acceptable price you choose.
All prices between CHF 0 and CHF 40 are equally likely to be selected as the randomly drawn offer price.
If SOLD: you will receive the randomly drawn offer price in exchange for the use indicated above.

(b) Valuations (Stage 2)

Figure 2: Screenshots of the interface

At the end of the experiment, the computer would randomly select one of the 40 report-viewing combinations from Stage 2 to be the one that could be implemented and determine subjects' earnings. Each combination could be selected with equal probability. The Becker-DeGroot-Marschak mechanism (Becker et al., 1964) was then employed to determine whether the subject received a randomly determined price between 0 and 40 CHF and the assessment reports were shown to the corresponding number of viewers, or whether the subject received no extra payments and no assessment report would be viewed by strangers. That is, the computer randomly selected a random price, p, between CHF 0 and CHF 40 and one of the combinations presented to the subject at Stage 2. Depending on whether p was higher than the indicated lowest acceptable price by the subject for that decision problem, two possible cases could then arise:

Case 1: If p was higher than or equal to the lowest acceptable price, the subject received the extra payment p for selling their personal information. The randomly drawn offer price was added to the payment, and the personal information was used in the manner specified in the report-viewing combination, meaning that each of the two reports would be shown to the specified number of viewers.

Case 2: If p was lower than the lowest acceptable price, the subject did not sell the personal information, and the randomly drawn offer price was not added to the payment for this experiment. In this case, the subject shredded all the personal data in the envelopes at the end of the experiment.

3.5 Stage 3: Incentivized Questions on Daily Privacy Behavior

In this stage, subjects answered incentivized questions about their daily-life privacy behaviors in various real-world contexts. Specifically, we collected measures of the amount of personal information shared with the public on Facebook, the number of active social media accounts, the number of store memberships or club programs in use, the settings of location services on the subject's mobile phone and the number of friends on Facebook.

For each of the above topics, subjects earned a basic amount of 1 CHF for providing an answer to that question and a potential bonus of 0.5 CHF for each additional piece of verifiable information they provided. For instance, when a subject was asked to indicate whether he or she shared different information items on Facebook, the subject could earn 0.5 CHF when indicating that he or she actually shared or did not share an item when the answers were verified to be true, but not when he or she indicated "I don't know." At the end of the experiment, each subject had a probability of 10 percent of having a random check in which the experimenter attempted to verify their answers. If *all* the information provided in Stage 3 was verified to be accurate, a subject received the potential bonus payments for providing accurate, verifiable information. If any part of a subject's responses were found to be inaccurate, all his or her payments from this Stage (ranging from 5 to 20 CHF if answered honestly) were forfeited. Hence, a subject unsure about an answer was instructed that it was in his or her best interest to honestly report "I don't know." Potential payoffs in each decision problem are shown to subjects instantaneously according to the chosen answers to make this incentive-compatible mechanism intuitive (see Appendix C for an example of the Stage 3 interface).

3.6 Stage 4: Questionnaire

Finally, subjects completed a questionnaire in Stage 4. This questionnaire included several demographic measures. Subjects were also asked their subjective satisfaction regarding their intelligence and body composition scores. Finally, we also elicited a single question measuring their general privacy attitude, using a variant of a question often used in the literature: "In general, how important you think privacy is?"

3.7 Sample

We recruited 225 subjects via the online registration tool h-root (Bock, Nicklish and Baetge, 2012) by sending e-mail invitations to registered participants from the general student population at the University of Zurich (UZH) and the Swiss Federal Institute of Technology in Zurich (ETH). The experiment was conducted at the Laboratory for Behavioral and Experimental Economics at UZH, on a rolling basis, with one subject at a time, to ensure subjects had full privacy control over their personal data even in the assessment stage. We implemented the experiment using the software z-tree (Fischbacher, 2007) and MATLAB.

Subjects were not informed of the nature of the experiment before arriving. However, upon arrival, they were given a detailed description of the procedures of the experiment and given the option to participate in an abbreviated experiment that offered a standard payment of CHF 20 and required no provision of sensitive personal information. This abbreviated experiment enables subjects to skip the assessment process, Stages 1 and 2, and to only participate in Stages 3 and 4. Compared to letting highly privacy-concerned subjects drop out directly, involving them in the abbreviated experiment provides us with information on their preferences. Out of 225 subjects, 7 subjects decided to participate only in the abbreviated experiment. There were also 2 subjects who decided not to participate at all after reading the written instructions.[11] Thus, in total, 216 subjects finished the experiment.

3.8 Follow-up Experiment Using Monetary Choices

In order to provide a benchmark comparison for the degree of rationality in our privacy experiment, we conducted a follow-up study using the same procedures and combination of budget sets in Stage 1, but instead of making choices over levels of personal information sharing, participants made choices from budget sets over monetary losses. Specifically, we translated the possible "loss" in privacy (corresponding to more viewers of one's personal information) to possible monetary losses by asking subjects to distribute losses between two accounts: a BLUE and a RED account. Subjects were endowed with a starting balance of 200 points in both accounts and knew that the exchange rate was 5 points = CHF 1. They then distributed possible losses, in points, between the two accounts in each of the 20 (4+16) decision problems in Stage 1 according to the downward-sloping budget sets in Figure 1(b). Following that, subjects went through Stages 3 and 4.[12] In total, 94 individuals participated, recruited from the same population and in the same manner as for our privacy experiment. The instructions and detailed analysis for this follow-up study are available in Appendix G and Appendix H.

4 Results

4.1 Rationality when Deciding Privacy Levels

4.1.1 Classification of privacy attitudes

Testing whether an individual's privacy choices are in conformance with GARP requires knowledge of the individual's privacy attitude toward the corresponding personal information. Privacy is recognized as a good for that information if one desires as few viewers as possible, and it is regarded as a bad if one wants as many viewers for that information as possible. An individual's privacy attitude, therefore, determines in which direction one's utility increases along with the privacy level, and GARP violations can be defined accordingly.

[11] The experiment was approved by the Ethics Commission of the Faculty of Business, Economics, Finance and Informatics at the University of Zurich.

[12] Stage 2 was excluded since eliciting willingness to accept for monetary losses does not make much sense. Some survey questions in Stage 4 were amended accordingly. For example, questions about subjective satisfaction regarding intelligence and body composition scores in the privacy experiment were replaced by questions about subjective preference for the color RED and BLUE.

We classify each of our subjects into one of four types of privacy attitudes: private (preferring less information sharing in both domains), public (preferring more information sharing in both domains) and two item-dependent types (keeping the body-composition report private while showing the intelligence report (IQ+) and keeping the intelligence report private while showing the body-composition report (BF+)). For this classification, we use subjects' choices in the unrestricted two-dimensional privacy-attitude questions from Stage 1A. We classify each of our subjects into the privacy attitude that generates predictions that are closest to the subject's choices in the two-dimensional privacy attitude questions.[13] We adopt this classification in the remainder of this study.

Subjects in our experiment exhibit heterogeneity of privacy attitudes. Figure 3(a) shows the percentage of subjects with different privacy attitudes, classified as described above, and the proportions of each type who exhibit consistency with GARP. In our experiment, 68 percent of subjects are private, and the remainder consists of around 14 percent of the IQ+ type, 8 percent of the BF+ type and 10 percent of the public type. Thus, most of our subjects exhibit a preference for sharing less personal information in this context, which is consistent with our expectations.

4.1.2 Conformity with GARP

Our first result concerns whether individuals willfully conceal and reveal personal information in a manner that is consistent with the General Axiom of Revealed Preference (GARP). Specifically, we apply GARP to our "report-viewing combination" data from Stage 1 by comparing how the allocation of privacy levels across the two information items varies according to the relative price of privacy—the relative numbers of viewers that have to be exchanged across the different information domains. Following Afriat (1967), if our finite dataset of report-viewing combinations generated by an individual is consistent with GARP, then the privacy-allocation choices can be rationalized by a utility function, U(x,y), for each of our subjects, with x and y being the number of viewers for different information items. We,

[13] We use 1-norm distance and define the closest privacy attitude as the attitude that minimizes the sum of the distances between the choices of the two-dimensional privacy attitude questions and the predictions for that privacy attitude. A stricter privacy-attitude classification is to require subjects classified as having a specific privacy attitude to choose within 10 percent of the unique point prediction for that attitude for both reports in both two-dimensional problems (within 1 percent of the whole choice set). This strict criterion results in 100 subjects remaining unclassified. The results using this classification are shown in Figure A.1.

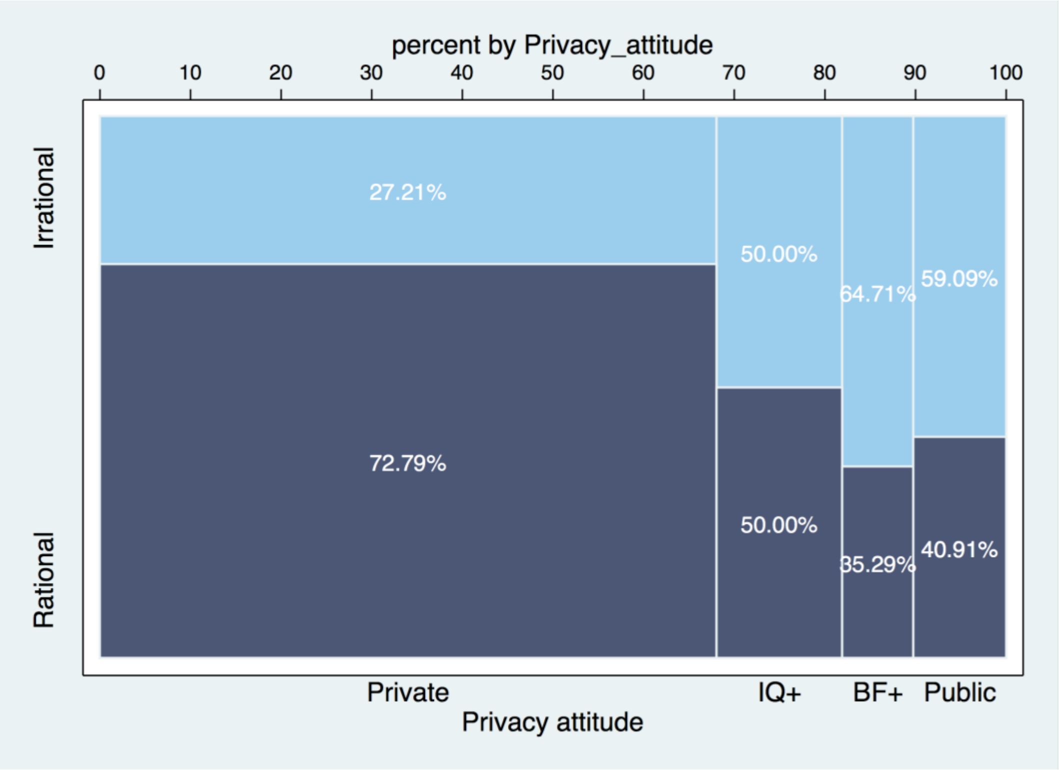

(a) Percentage of choices satisfying GARP by privacy type

Notes: The plot depicts the proportion of rational and irrational subjects within each privacy attitude. The width of the bar indicates the percentage of subjects of the corresponding privacy attitude.

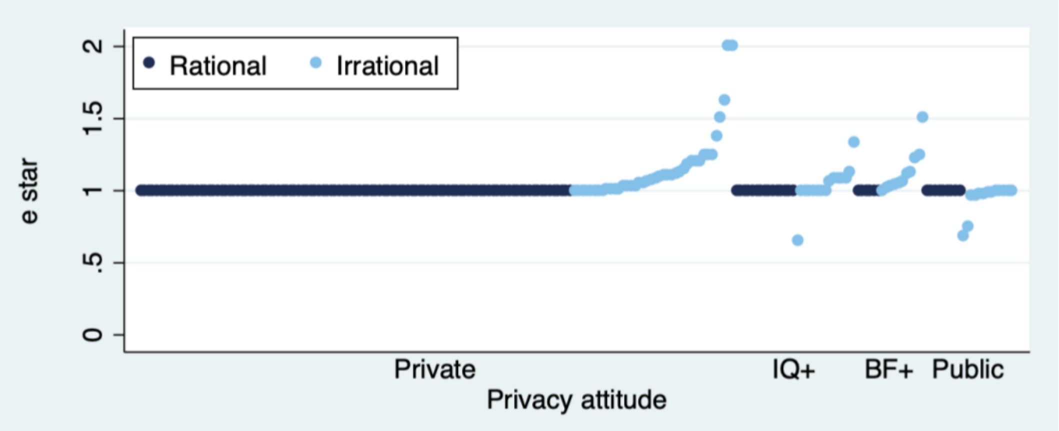

(b) Efficiency score: CCEI (e^*)

Notes: The scatterplot depicts the distribution of the efficiency measure, e^*, for subjects with different privacy attitudes. Each dot represents the e^* of one subject. Within each attitude, rational subjects (subjects with no GARP violations) are plotted first with navy dots. The e^* of irrational subjects are sorted in ascending order, represented by light blue dots.

Figure 3: Rationality measures for Stage 1 privacy choices

therefore, test whether individuals allocate privacy levels across items in a consistent and consequently rational manner that can be modeled using utility functions.[15]

Our primary interest in Stage 1 is to identify the degree to which individuals' privacy choices exhibit consistency with GARP. To measure the degree of rationality for each type, we adapt the rationality test according to the different privacy attitudes—i.e., the direction of preferences on the two dimensions of information sharing. Using these measures, we find in total, more than 63 percent of subjects exhibit perfect consistency with GARP. As shown in Figure 3(a), the degree of rationality is the highest among those individuals that we separately classify as private types: more than 72 percent of the private subjects are rational. The rational percentages corresponding to the IQ+, BF+ and the public types are 50 percent, 35.29 percent, and 40.91 percent, respectively.[17]

Another way to look at the degree of rationality of privacy decisions is to measure how closely individual privacy choices adhere to GARP. It is possible that GARP violators are choosing in a manner that is very close to satisfying GARP, and the cost of the violations is minimal. In this case, we may not want to conclude that they fail this standard of economic rationality when making privacy choices. We address this concern by adapting the Critical Cost Efficiency Index (CCEI) efficiency measure e^* in Afriat (1972) to different types of privacy attitudes. If a subject values privacy as a good (bad), then $|1\text{-}e^*|$ indicates the privacy level loss (gain) by the proportion of the increased (decreased) number of viewers needed to remove all GARP violations for that individual.

Figure 3(b) shows the e^* scores, arranged in increasing order, for individuals with different privacy attitudes. The average e^* scores for each privacy type are also displayed in Table 2. This score averaged 1.0425 for the private types and 0.9670 for the public types, which implies that, on average, the privacy level needs to be reduced by about 4 percent to eliminate all GARP violations by a private subject and increased by about 3 percent to rationalize the choices of a public subject.

To provide some context for these numbers, Table 2 compares the percentage of subjects acting consistently with GARP in this study with several previous experimental studies

[15] According to Varian (1982), GARP requires that if a report-viewing combination C is revealed preferred to C', then C' is not strictly and directly revealed preferred to C. That is, the combination C must require a privacy level at least as high (low) as C' at the relative prices where C' is chosen if one considers privacy a good (bad) for the corresponding personal information.

[17] We do not find evidence that irrational types and public types spend less time making their choices in this study. Specifically, we find no difference in the response time of subjects for the choices in Stage 1 when comparing either rational vs. irrational types or public vs. private types.

testing consistency with GARP when making other types of tradeoffs—between consumption goods (Cox, 1997; Sippel, 1997; Mattei, 2000; Harbaugh et al., 2001), monetary payoffs between one's self and another (Andreoni and Miller, 2002, Fisman et al., 2007) and between two possible monetary payoffs for one's self (Choi et al., 2007, Choi et al., 2014). The table also reports the Bronars' power, a measure of how likely the budget sets employed are to detect violations by an individual choosing entirely at random. The Bronars power in previous studies ranges from 0.243 to 1, with recent studies using high power levels, close to or equal to 1. Our power for private types who value privacy as an economic good is sufficiently high (0.9996) to be comparable to these studies.[18] For public types, our Bronars power (0.9419) is slightly lower than in recent studies.

Table 2: Comparison of degree of rationality across different studies

	Tradeoffs	Bronars Power	Rational %	CCEI
Cox, 1997	Consumption goods	0.243-0.664	63.2	
Sippel, 1997	Consumption goods	0.613-0.973	36.8-58.3	
Mattei, 2000	Consumption goods	0.989	68	
Harbaugh et al., 2001	Consumption goods	0.98	65	0.94
Andreoni and Miller, 2002	Money	0.781-0.947	89.8	
Choi et al., 2007	Money	1	(64.6)[a]	0.954
Fisman et al., 2007	Money	1	10.5[b] (53.9)	
Choi et al., 2014	Money	1	22.8 (45.2)	0.881
This paper	Privacy	Private: 0.9996	72.8 (83.7)	1.0425
		Public: 0.9419	40.9 (90.9)	0.9670
	Monetary losses[c]	0.9996	45.1 (78.0)	1.0599

"Rational %" generally refers to the percentage of subjects exhibiting no GARP violations.
[a] The numbers in parentheses classify as rational any subject with CCEI scores above 0.95. This threshold is adapted to 1.05 for private types in our privacy study and for subjects in our follow-up study with monetary losses.
[b] This number is not reported in Fisman et al. (2007) but instead in their working-paper version (2005). The low number potentially reflects numerous small violations, perhaps due to the choice interface they employ.
[c] For our study with monetary losses, we only include subjects classified as valuing money (91 of 94).

[18] The power in our experiment is a design choice: the three-stage design of this experiment and the feature that Stage 2 needs to double the number of choices of Stage 1 both limit the number of decisions in our GARP test. Moreover, given the number of decisions, shifting budget lines to increase the power of private types will decrease the power for public types, since their utility increases in opposite directions. Hence the power of both types in our experiment is a result of tradeoffs given that we expect more, and are primarily interested in, private types.

Among private types, the proportion of rational subjects in our privacy study (ranging from 72.8 to 83.7 percent, depending on whether one uses the 0.95 CCEI threshold proposed by Varian (1991) to allow for minor violations) is higher than that in most of the previous studies, except in Andreoni and Miller (2002),[20] and the percentage of rational public subjects ranges from 40.9 to 90.9 percent, which is similar to the range observed in other studies (36.75 to 89.8 percent) and is higher than most recent studies when using the 0.95 CCEI threshold. Hence, overall, the proportion of individuals who allocate privacy levels rationally in our study does not seem to differ substantially from the rational proportion of individuals facing other kinds of economic tradeoffs in earlier studies.[21]

Of course, GARP violations and the CCEI index may be influenced by the specific combination of budget sets used across different studies, making the comparisons in Table 2 hard to interpret. Therefore, we also conducted a follow-up condition in which a separate group of participants recruited from the same population as those in our main privacy experiment made identical choices to those in Stages 2 and 3, but with respect to combinations of monetary losses (for details of the design and detailed analysis, see Appendix G and Appendix H). As in our main experiment, we used an initial set of unrestricted two-dimensional choices to classify the orientation of these participants' preferences toward money. Not surprisingly, a very large majority (97%) of participants are classified as valuing money, meaning that they selected bundles that minimized their total monetary losses. The bottom row of Table 2 summarizes the degree of rationality observed in the choices made by such money-valuing types in this follow-up experiment. The average CCEI of the types who like money (1.0599) is not statistically different from the CCEI for the private types (1.0425) making privacy tradeoffs. Therefore, the perturbation that is required for rationalizing privacy choices is similar when compared to the degree that is necessary to eliminate the GARP violations when making monetary decisions or decisions between consumption goods. This comparison shows that people making tradeoffs involving privacy exhibit similar consistency as individuals making monetary tradeoffs.

4.2 Consistency between Monetary Equivalents and Privacy Choices

[20] The difference with Andreoni and Miller is unsurprising, as our subjects were facing a more stringent test that provides more opportunities to violate GARP. This is reflected in a comparison of the Bronars power.

[21] Notice that here we are comparing tradeoffs that are made "within" the same domain (money versus money, privacy versus privacy, etc.). This type of tradeoff tells us how consistent an individual is when making choices in that domain, but it does not tell us how consistent an individual will be when making cross-domain choices (e.g., trading off privacy for money). In Section 4.2, we will consider this type of tradeoff.

The second part of our results informs the possibility of efficiently allocating privacy by monetizing personal information. This part extends the study of the rationality of privacy choices to cases involving tradeoffs between monetary benefits and privacy protection. It investigates the consistency between both the preference relations and privacy attitudes revealed in Stage 1 and their monetary equivalents.

4.2.1 Consistency between Monetary Equivalents and the Privacy Preference Relation

First, we construct a measure of efficiency loss when individuals translate their privacy preferences into monetary valuations. We use the directly revealed preference relation in each decision problem involving a downward-sloping budget set from Stage 1—i.e., since, for each budget line from Stage 1, one bundle presented in Stage 2 was chosen and the other was not—and compare this relation with the monetary equivalents for the chosen and unchosen options from Stage 2. To be specific, in Stage 2, we elicit the willingness-to-accept (WTA) for the 16 choices selected by a subject and for the 16 perturbed choices selected by the computer. Since the former is directly revealed preferred to the latter, when the monetary equivalent of a perturbed choice is smaller than that of the chosen choice on the same budget line, the difference between the two WTAs signals the size of efficiency loss when people are asked to represent their privacy preference in monetary values. We use the sum of the differences across all 16 budget lines, which we refer to as the monetary-value noise (MVN), as our measure of efficiency loss in the translation of directly elicited privacy preferences to monetary valuations.[22]

Overall, 46.3 percent of subjects display perfect consistency (i.e., MVN = 0) when representing their choices from Stage 1 in terms of monetary equivalents. However, the remaining 53.7 percent of subjects exhibit some degree of inconsistency when indicating their privacy preferences in money, and these inconsistent subjects include both rational and irrational types. To be specific, 47 percent of rational subjects and 65 percent of irrational subjects show preference reversals when representing their Stage 1 preference in monetary terms in Stage 2. This result indicates that a significant proportion of subjects exhibit privacy preference reversals when assessing privacy preference with monetary equivalents, even for

[22] Recall that in Stage 2, subjects provided monetary valuations for 40 bundles. Of these, 16 were the bundles selected by subjects from the 16 downward-sloping budget sets and another 16 were different bundles selected from the same budget sets by the computer. The remaining 8 bundles consisted of the 2 selected bundles from the one-dimensional choices, the 2 selected bundles from the two independent two-dimensional choices subjects made at the beginning of Part 1, and 4 pre-determined bundles used to elicit valuations that are comparable across all subjects. Our primary focus in this section is on the valuations for the 16 matched pairs of bundles from the same Part 1 budget sets.

subjects who seem to rationally decide the privacy levels for different personal information items when monetary tradeoffs are not involved.

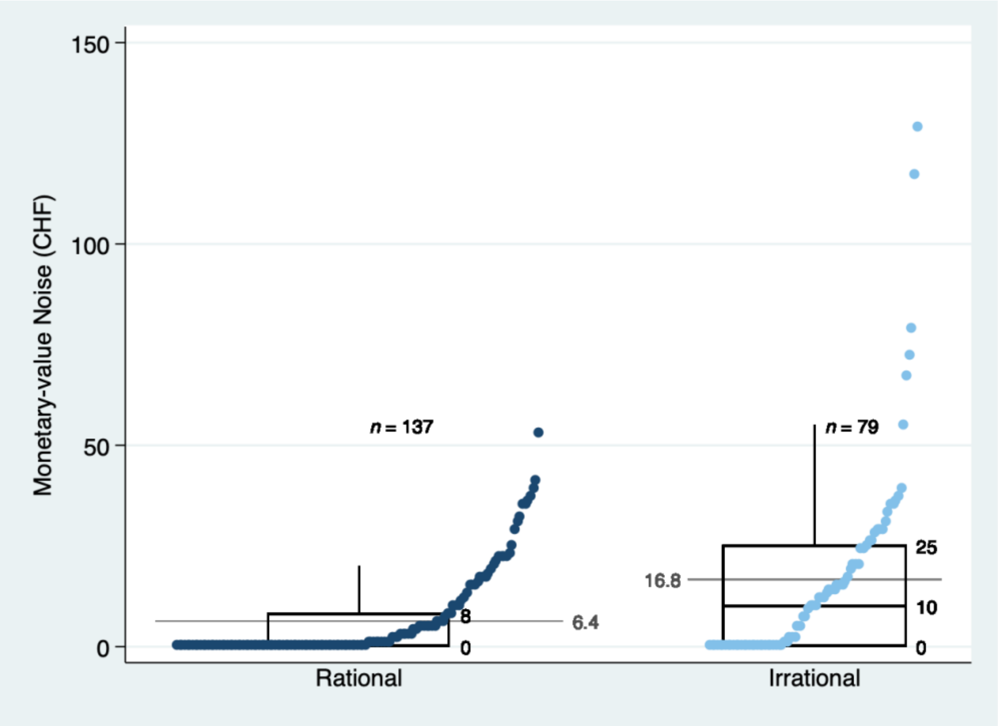

Notes: The box plot depicts the quartiles of monetary-value noise for rational and irrational subjects. Within each type, the dots represent individuals' monetary-value noise sorted in ascending order; mean of monetary-value noise is indicated next to the gray horizontal line.

Figure 4: Monetary-Value Noise (Stage 2)

Despite the fact that about half of both the rational and irrational privacy types exhibit noise when representing their preferences in monetary units, there is a stark difference between the size of noise between these two groups. Figure 4 shows that the average total MVN (the grey bar) across 32 choices for those subjects classified as irrational is about CHF 16.76, which is more than 2.6 times of the size of the average total noise of rational subjects: CHF 6.38.[23] Since the standard deviation of MVN is smaller for rational subjects than for irrational subjects: we apply a t-test for unequal variance data and show that this difference is significant at the 0.1% level (t = −3.572). Notice that this difference is not mechanical; GARP violations in Part 1 do not preclude a subject from having zero MVN in Part 2, since the MVN is always computed using the monetary valuations of a chosen bundle and a similar unchosen one.

[23] Figure A.2 in the Appendix presents the proportion of subjects who exhibit the inconsistency for each preference category. Figure A.3 in the Appendix shows the cumulative distribution functions of MVN for the irrational types and rational types by privacy attitudes. For all four types of privacy attitude, irrational individuals have larger MVN than rational ones.

Moreover, it is not driven by subjects whose behavior makes them outliers, since the MVN of irrational types is higher than rational types for each quartile. Therefore, the comparisons show that trading privacy against money may induce very different monetary consequences for people who do and do not obey GARP in the simple tradeoff tasks in Stage 1.[24]

4.2.2 Consistency between Monetary Equivalents and Privacy Attitudes

As Figure 4 shows, the size of monetary-value noise of privacy varies substantially across individuals. A natural question to ask is whether monetary equivalents are representative of the underlining privacy attitudes. The answer is important because monetary equivalents are extensively used for accessing the value of privacy.

Our results confirm that monetary equivalents can discriminate between the underlying privacy attitudes of different types of individuals. We compare the average monetary compensation required by private and public types for the four preset bundles that hold fixed, across participants, the combinations of viewers: (20,20), (65,65), (0,65) and (65,0). For these combinations, we have WTA values for identical bundles for every participant in our study. We then calculate $\text{WTA}_{\text{preset}}$, the averaged WTAs of these four preset combinations, for each of our subjects and compare the mean of the $\text{WTA}_{\text{preset}}$ for private and public subjects, according to the classification in Stage 1.

Figure 5 presents the averaged $\text{WTA}_{\text{preset}}$ for each subject, as well as the quartiles and means for private and public types. It shows that the demanded price for the private types tends to be higher than that of the public types at every quartile. The mean $\text{WTA}_{\text{preset}}$ (the grey line on the box plot) of private subjects (18.61 Swiss francs) is significantly higher than that of public subjects: 11.35 francs, confirmed by a t-test with unequal variance at the 5% level (t = 2.6608).[25] Hence, despite the heterogeneity and the noise present in individuals' privacy choices involving monetary tradeoffs, the results in this part suggest that there is value in employing monetary equivalents as a privacy attitude measure, at least at the group level.

[24] Table A.2 in the Appendix shows that a larger MVN correlates with a lower efficiency level (measured by CCEI) for both private and public types after controlling for personal characteristics, although the result is not statistically significant for public types. For private types, the size of monetary noise is significantly larger (t = 2.35) for individuals who had greater inconsistency with GARP when trading off privacy between two information items (i.e., a larger e^*). For public types, the relation is not significant (t = -0.11), although the size of monetary noise is larger when $e^*$is smaller. We find no difference between the response times of public and private types when they are trading off privacy control for money (Stage 2).

[25] Figure A.4 in Appendix further shows that this trend holds for both the irrational and the rational types.

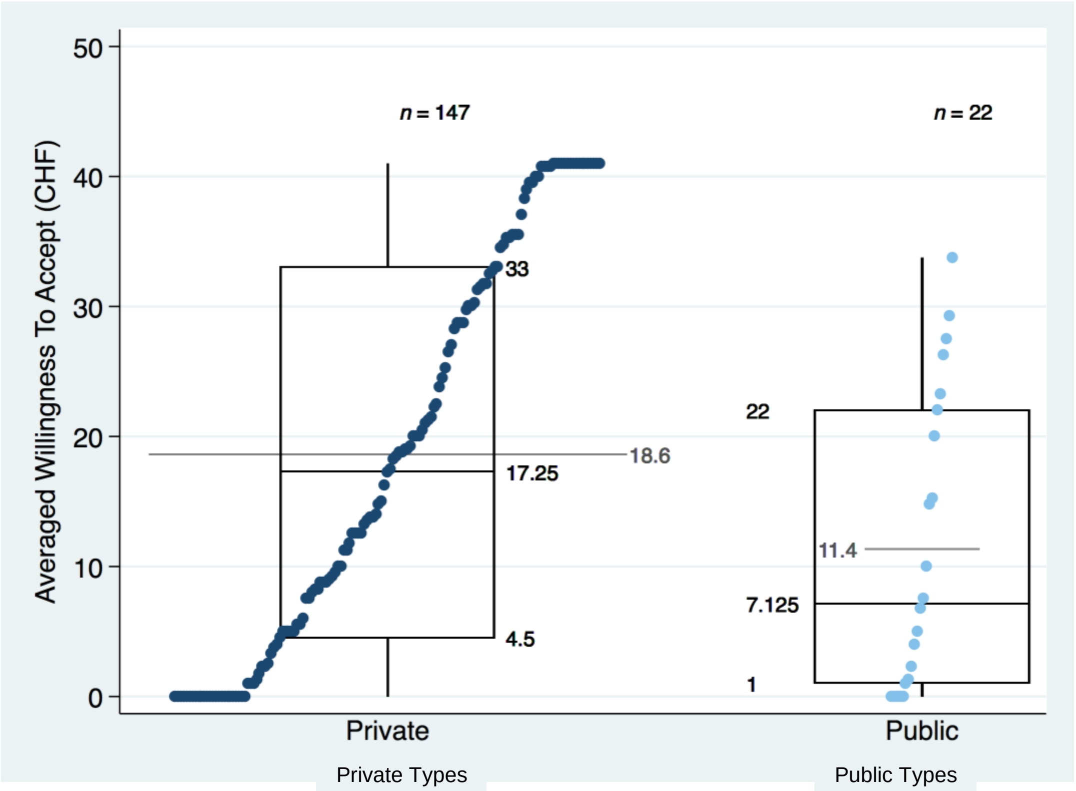

Notes: The box plot depicts the quartiles of the willingness-to-accept of sharing personal information according to the preset bundles for subjects of private and public types. Within each type, the dots represent individuals' willingness-to-accept sorted in ascending order; mean of willingness-to-accept is indicated next to the gray horizontal line.

Figure 5: Valuations for WTA$_{preset}$ by privacy attitude (private and public types)

4.3 Do laboratory preferences correlate with attitudes and behaviors?

The third part of our analysis investigates whether subjects' privacy attitudes and their degree of rationality in our laboratory setting exhibit some relationship with subjects' daily-life privacy behavior outside the laboratory. The existence of a correlation is necessary but not sufficient for our measures to reliably proxy for more general forms of real-life behavior. Hence, we cautiously aim here only to conduct an exploratory test the relationship between our lab measures and some real-world behaviors.

4.3.1 Measures of Daily-life Information Sharing

To test whether privacy attitudes and measures of rationality in this study correlate with how people behave in more natural privacy decisions, we utilize the incentivized daily-life information sharing behavior data from Stage 3 to construct four variables: (1) *Facebook Public*: the number of information items a subject shares with all Facebook users on their profile page,[27] including individuals not on the friend list; (2) *Social Media Public*: the number of social networking accounts on which a subject posts public content that anyone using the media can see;[28] (3) *Memberships*: the number of categorized memberships or club programs a subject uses;[29] and (4) *Location shared*: whether one activated the location services on one's smartphone.[30] Individuals who chose to leave one of these questions unanswered are excluded from the analysis in this section, which leaves us 198 out of 216 subjects.[31]

We first conduct an exploratory principal-component factor analysis using the four variables of real-world data sharing behaviors as inputs, with orthogonal Varimax rotation, and obtain two factors of daily-life information sharing behavior: sharing in social networks and the exchange of personal information for money or services.[32] We then investigate the extent to which these two factors correlate with our individual-level privacy-preference measures. Specifically, we study whether privacy attitudes or the rationality of privacy preferences are predictive for heterogeneity in real-life information sharing at the individual level, controlling for socioeconomic and personal characteristics that have been documented to be influential for personal data sharing decisions.

Table 3 presents ordinary-least-squares regression results, using either the measure of sharing in social networks (factor 1) or the exchange of personal information for money and services (factor 2) as the dependent variable. The explanatory variables of interest are indicators for whether a participant is classified as a private type or as an irrational type based on his or her choices in Stage 1. Model 1 shows that individuals who are public types in our experiment share significantly more personal data in social networks than private types (by 0.91 standard deviation of factor 1). Model 2 confirms that this result is robust after controlling for several important socioeconomic and personal characteristics in the privacy literature:

[27] We consider eight items on the overview page of a Facebook user, including the profile picture, workplace, educational background, phone number, address, e-mail address and month, day and year of birth. If a subject has no Facebook account, we consider this as zero items shared. If a subject has a Facebook account but does not know what they have shared, we use the default Facebook setting that three items are public.

[28] The seven forms of social media we consider are: Facebook, Google+, Instagram, LinkedIn, Pinterest, Twitter and YouTube.

[29] Six categories of membership are considered: the first two categories are two very popular retail stores in Switzerland: coop and MIGROS; the other four store categories are furniture/household (IKEA, interio, Jumbo, Pfister, etc.), clothing (C&A, ESPRIT, H&M, PKZ, etc.), electronics (Digitec, InterDiscount, Media Markt, etc.) and Online shopping platforms (Amazon, GALAXUS, zalando, etc.).

[30] If a subject indicates that he or she uses a smartphone and doesn't know if the locational service is on, we consider the location service activated, because this is the default option for most smartphones.

[31] Allowing individuals to indicate that they would like to leave a question unanswered enables us to distinguish people who do not want to answer the question from those who don't know the answer and those who do not use the service or accounts. Please see Section 3.5 for details of the design.

[32] These two factors yield eigenvalues greater than 1 (see Appendix D for details of the factor analysis).

income, sex, education and marital status. Models 3 and 4 suggest that rationality has no effect on sharing in social networks. Models 5 and 6 show that privacy attitudes have no impact on the exchange of personal information for money and services. However, Models 7 and 8 provide modest evidence that that irrational types tend to trade their personal information for money or services more often than rational types (by 0.25 standard deviation of factor 2), though the results are only significant at the 10 percent level.

Table 3: Relationships between privacy-preference measures and daily-life data sharing

	Sharing in social networks (Factor 1)				Exchange of personal information for money and services (Factor 2)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Public	0.910***	0.859***			−0.078	−0.004		
	(4.02)	(3.73)			(−0.32)	(−0.02)		
Irrational			0.17	0.208			0.252+	0.250+
			(1.15)	(1.41)			(1.72)	(1.66)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Constant	−0.178*	0.131	−0.062	−0.017	0.019	0.061	−0.092	−0.496
	(−2.17)	(0.35)	(−0.70)	(−0.05)	(0.21)	(0.14)	(−1.04)	(−1.41)
N	152	152	198	198	152	152	198	198

Notes: Controls include income (categorical), sex (dummy), education (categorical) and marital status (categorical). The two item-dependent types (IQ+ and BF+) are omitted for regressions using Public as an independent variable (i.e. models (1), (2), (5) and (6)). Results are robust when including item-dependent types (coded as either public=0 or public=1) in these models; see Table A.1 for more details. t statistics in parentheses, + p < .10, * p < .05, ** p < .01, *** p < .001.

4.3.2 General Privacy Attitude and Privacy Choices

Our final analysis examines the consistency between individuals' self-reported general privacy attitudes and their behavior in the experiment. People who claim to care about privacy seldom demonstrate comparable concerns in their daily-life data sharing behavior. This discrepancy between the stated privacy attitudes and behaviors is known as the "privacy paradox." (Athey et al., 2017; Kokolakis, 2017). We test if this inconsistency persists even in the very simple setting that we study.

Our results indicate a high consistency between the choice-based privacy attitudes measured in the lab and the self-reported general privacy attitude, elicited using a widely employed survey question, for both the irrational and the rational types. The General Privacy Attitude is measured by an answer on 5-point Likert scale, for the question "In general, how important do you think privacy is?”[33] Table 6 reports ordered-logistic regression results, using the General Privacy Attitude as the dependent variable; the outputs are proportional odds ratios. Model 1 shows that the individuals who are classified as public types based on their Stage 1 choices in our experiment report significantly lower general privacy concerns (z= -4.11) than private types. Model 2 confirms that this result is robust after controlling for several important socioeconomic and personal characteristics in the privacy literature: income, sex, education and marital status. Therefore, our choice-based classification of privacy attitudes appears consistent with the self-reported attitude that is frequently used in research and policy.

Table 6: Correlation between Choice-based Privacy Attitudes and Self-reported General Privacy Attitudes

	General Privacy Attitude	
	(1)	(2)
Public	0.1342^{***} (-4.11)	0.1317^{***} (-3.94)
Controls	No	Yes
/cut1	-4.2753	-4.3266
/cut2	-1.0438	-1.0139
/cut3	1.7083	1.8017
N	169	169

Notes: Ordered logistic regression; coefficients are proportional odds ratios. Controls include income(categorical), sex (dummy), education (categorical) and marital status (categorical). The two item-dependent types (IQ+ and BF+) are omitted since we use Public as an independent variable. z statistics in parentheses, + p < .10, * p < .05, ** p < .01, *** p < .001

5 Conclusion

This paper provides evidence on people’s ability to rationally manage the sharing of their personal data. We study this by presenting subjects with decisions involving tradeoffs between sharing different kinds of personal information items and observing the extent to which their choices satisfy basic standards of economic rationality.

More precisely, we present a novel choice experiment consisting of three main parts. First, people confront decisions in which they allocate privacy levels between different personal information items. Second, they decide at what prices they are willing to trade their personal data for money for various uses. At the end of the experiment, we also collect incentivized details of daily-life information sharing and survey subjects about their general

[33] We elicit the answer at the end of the experiment as one of the questions in Stage 4. Subjects choose from five degrees of importance: Unimportant, Somewhat unimportant, Somewhat important, Very important, and Critical.

privacy attitudes. Through these three kinds of data, we provide several insights into the rationality and consistency of privacy behavior.

First, we document considerable heterogeneity in both privacy attitudes in the domains in our experiment and in the rationality with which people confront tradeoffs in these domains. A majority of individuals in our data want to minimize the information shared, while others desire to share one or both pieces of information. Despite this heterogeneity of privacy attitudes, most individuals allocate privacy levels across different information domains in a manner that is consistent with the Generalized Axiom of Revealed Preference. Thus, there seems to be at least the possibility that people can manage simple tradeoffs sensibly when deciding how to balance different dimensions of personal information sharing. Most people in our sample are capable of sharing or concealing personal information as if they are maximizing an underlying preference ordering in a manner consistent with utility maximization. However, it is important to note that we study a very simple choice context in which there are only two domains of information sharing, the consequences of sharing public data are simple and known, and we make it easy for people to exert full control over their personal data. The fact that a substantial proportion of subjects exhibit violations of even very simple notions of rationality in this setting suggest that rationality may be much lower as one moves to more complex natural settings. Thus, our observation that roughly two-thirds of individuals exhibit high levels of rationality presents more of an upper bound than a number likely to generalize to other contexts. Nevertheless, it is worth noting that the degrees of rationality we observe in the privacy domain are comparable to those observed in other domains, both in other studies and in a novel follow-up study we conducted.

Second, there is further inconsistency when translating privacy preferences into monetary values. While a large proportion of subjects exhibit consistency in the tradeoffs they make between sharing of different personal information items and the monetary valuation of their personal information, there are also many subjects whose choices exhibit preference reversals. More than half of our subjects exhibit such reversals. Moreover, the people who are irrational when allocating privacy levels to different information items forgo a much larger amount of money (2.6 times as much) through their inconsistency than do rational types. These monetary welfare effects, therefore, cast some doubt on the feasibility of improving the efficiency of privacy policy by monetizing or licensing personal information, at least for a substantial subset of the population. Indeed, only a relatively small proportion—less than a third—of subjects in our experiment demonstrate perfect consistency across all choices in Stage 1 and Stage 2.

Finally, we also provide exploratory evidence on the extent to which privacy preferences in different domains are connected. First, monetary valuations, despite the noise at the individual level, nevertheless capture the underlying privacy preferences of groups of private and public types. Second, individuals who are irrational when deciding privacy levels for different information items squander larger amounts when trading off privacy control for money. Third, laboratory privacy preference measures appear to have some relationship to at least some domains of real-world privacy behaviors and to general self-reported privacy attitudes, measured with a question that is often used in research and policy. Therefore, our behavioral measures not only provide insights into the internal consistency of privacy preferences in the laboratory, but are also informative about behaviors and attitudes outside the laboratory.

Overall, our results provide novel evidence on three fundamental privacy issues: (i) the extent to which people manage their personal information rationally; (ii) the feasibility of improving the efficiency of privacy rights by liberalizing the monetizing and licensing of personal information; and (iii) the heterogeneity in privacy attitudes and the degree of rationality of privacy preferences and their implications for information sharing in a variety of contexts encountered in modern life. A better understanding of these issues is critical for the design of effective policies that balance the benefits of information sharing against potential costs from individuals' irrationality.

References

Abowd, J. M., & Schmutte, I. M. (2019), ‘An economic analysis of privacy protection and statistical accuracy as social choices’, American Economic Review, 109(1), 171-202.

Acemoglu, D., Makhdoumi, A., Malekian, A., & Ozdaglar, A. (2021), ‘Too Much Data: Prices and Inefficiencies in Data Markets. Forthcoming in American Economic Journal: Microeconomics.

Acquisti, A., Brandimarte, L. and Loewenstein, G. (2015), ‘Privacy and Human Behavior in the Age of Information’, Science 347(6221), 509–514.

Acquisti, A., Brandimarte, L., & Loewenstein, G. (2020), ‘Secrets and likes: the drive for privacy and the difficulty of achieving it in the digital age’, Journal of Consumer Psychology, 30(4), 736-758.

Acquisti, A., John, L. K. and Loewenstein, G. (2013), ‘What Is Privacy Worth?’, The Journal of Legal Studies 42(2), 249–274.

Acquisti, A. and Ralph, G. (2006), ‘Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook’, Privacy Enhancing Technologies pp. 1–22.

Acquisti, A., Taylor, C. and Wagman, L. (2016), ‘The Economics of Privacy’, Journal of Economic Literature 54(2), 442–492.

Afriat, S. N. (1967), ‘The Construction of Utility Functions from Expenditure Data’, International Economic Review 8(1), 67–77.

Afriat, S. N. (1972), ‘Efficiency estimation of production functions’, International Economic Review 13(3), 568–598.

Athey, S., Catalini, C., & Tucker, C. (2017), ‘The digital privacy paradox: Small money, small costs, small talk (No. w23488)’, National Bureau of Economic Research.

Becker, G. M., Degroot, M. H. and Marschak, J. (1964), ‘Measuring utility by a single-response sequential method’, Systems Research and Behavioral Science 9(3), 226–232.

Benndorf, V. and Normann, H.-T. (2017), ‘The Willingness to Sell Personal Data’, Scandinavian Journal of Economics. Accepted Author Manuscript,. doi:10.1111/sjoe.12247

Beresford, A. R., Ku¨bler, D. and Preibusch, S. (2012), ‘Unwillingness to pay for privacy: A field experiment’, Economics Letters 117(1), 25–27.

Bock, O., Nicklisch, A. and Baetge, I. (2012), ‘ hroot: Hamburg registration and organization online tool’, H-Lab Working Paper (1).

Bronars, S. G. (1987), ‘The power of nonparametric tests of preference maximization’, Econometrica 55(3), 693.

Burtch, G., Ghose, A., & Wattal, S. (2015), ‘The hidden cost of accommodating crowdfunder privacy preferences: A randomized field experiment’, Management Science, 61(5), 949-962.

Carrascal, J.P., Riederer, C., Erramilli, V., Cherubini, M., de Oliveira, R. (2013), ‘Your browsing behavior for a big mac: Economics of personal information online’, Proceedings of the 22nd international conference on world wide web (pp. 189–200).

Casadesus-Masanell, R., & Hervas-Drane, A. (2015), ‘Competing with privacy’, Management Science, 61(1), 229-246.

Choi, J. P., Jeon, D. S., & Kim, B. C. (2019), ‘Privacy and personal data collection with information externalities’, Journal of Public Economics, 173, 113-124.

Choi, S., Kariv, S., Muller, W. M. and Silverman, D. (2014), ‘Who Is (More) Rational?’, American Economic Review 104(6), 1518–50.

Danezis, G., Lewis, S., Anderson, R.J. (2005), ‘How much is location privacy worth?’, Fourth workshop on the economics of information security.

Dwork, C. (2008), Differential Privacy: A Survey of Results, in ‘Theory and Applications of Models of Computation’, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 1–19.

European Parliament, Council of the European Union (2016), ‘REGULATION (EU) 2016/ 679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL - of April 27 2016 - on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/ 46/EC (General Data Protection Regulation)’, pp. 1–88.

Fischbacher, U. (2007), ‘z-Tree: Zurich toolbox for ready-made economic experiments’, Experimental Economics, 10(2), 171-178.

Ghosh, A. and Roth, A. (2015), ‘Selling privacy at auction’, Games and Economic Behavior 91(C), 334–346.

Goldfarb, A., & Tucker, C. (2011), ‘Online display advertising: Targeting and obtrusiveness’, Marketing Science, 30(3), 389-404.

Hirshleifer, J. (1971), ‘The Private and Social Value of Information and the Reward to Inventive Activity’, American Economic Review 61(4), 561–574.

Hirshleifer, J. (1980), ‘Privacy: Its Origin, Function, and Future’, The Journal of Legal Studies 9(4), 649–664.

Hoofnagle, C. J. and Urban, J. M. (2014), ‘Alan Westin’s Privacy Homo Economicus’, Wake Forest Law Review 49, 261–321.

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

Kalai, G., Rubinstein, A. and Spiegler, R. (2002), ‘Rationalizing Choice Functions by Multiple Rationales’, Econometrica 70(6), 2481–2488.

Kokolakis, S. (2017), ‘Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon’, Computers & Security 64(C), 122–134.

Kummer, M., & Schulte, P. (2019), ‘When private information settles the bill: Money and privacy in Google’s market for smartphone applications’, Management Science, 65(8), 3470-3494.

Laudon, K. C. (1996), ‘Markets and privacy’, Communications of the ACM 39(9), 92–104.

Miller, A. R., & Tucker, C. (2009), ‘Privacy protection and technology diffusion: The case of electronic medical records’, Management science, 55(7), 1077-1093.

Miller, A. R., & Tucker, C. (2018), ‘Privacy protection, personalized medicine, and genetic testing’, Management Science, 64(10), 4648-4668.

Nissenbaum, H. (2009), Privacy in Context, Technology, Policy, and the Integrity of Social Life, Stanford University Press.

Peukert, C., Bechtold, S., Batikas, M., & Kretschmer, T. (2020), ‘European privacy law and global markets for data’, Available at SSRN 3560392.

Posner, R. A. (1978), ‘The Right of Privacy’, Georgia Law Review.

Posner, R. A. (1981), ‘The Economics of Privacy’, American Economic Review 71(2), 405–409.

Samuelson, P. (2000), ‘Privacy As Intellectual Property?’, Stanford Law Review 52(5), 1125–1173.

Schudy, S. and Utikal, V. (2017), ‘You must not know about me’-On the willingness to share personal data’, Journal of Economic Behavior & Organization, 141: 1-13.

Schwartz, P. M. (2004), ‘Property, privacy, and personal data’, Harvard Law Review 117(7), 2056.

Sleeper, M., Balebako, R., Das, S., McConahy, A. L., Wiese, J., & Cranor, L. F. (2013), ‘The post that wasn't: exploring self-censorship on facebook’, In Proceedings of the 2013 conference on Computer supported cooperative work (pp. 793-802).

Solove, D. J. (2013), ‘Privacy Self-Management and the Consent Dilemma’, Harvard Law Review pp. 1880–1903.

Stigler, G. J. (1961), ‘The Economics of Information’, Journal of Political Economy 69(3), 213–225.

Stigler, G. J. (1962), ‘Information in the Labor Market’, Journal of Political Economy 70(5), 94–105.

Stigler, G. J. (1980), ‘An Introduction to Privacy in Economics and Politics’, The Journal of Legal Studies 9, 623–644.

Simonite, T. (2014), ‘Sell your personal data for $8 a month’, MIT Technology Review, 117(3), 20–20.

Spiekermann, S., Korunovska, J., Bauer, C. (2012), ‘Psychology of ownership and asset defense: Why people value their personal information beyond privac’, Available at SSRN 2148886.

Stutzman, F., Gross, R. and Acquisti, A. (2012), ‘Silent listeners: The evolution of privacy and disclosure on facebook’, Journal of privacy and Confidentiality 4(2), 7–41.

Taylor, C. R. (2004), ‘Consumer Privacy and the Market for Customer Informa- tion’, The RAND Journal of Economics 35(4), 631.

The Privacy Office, US Department of Homeland Security (2009), ‘The Fair Information Practice Principles: Framework for Privacy Policy at the Department of Homeland Security’, pp. 1–4.

Tsai, J. Y., Egelman, S., Cranor, L. and Acquisti, A. (2011), 'The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study', Information Systems Research 22(2), 254–268.

Varian, H. R. (1982), 'The Nonparametric Approach to Demand Analysis', Econometrica 50(4), 945.

Varian, H. R. (1991), 'Goodness-of-Fit for Revealed Preference Tests', Ann Arbor: Department of Economics, University of Michigan.

Varian, H. R. (2009), 'Economic Aspects of Personal Privacy', Internet Policy andEconomics (Chapter 7), 127–137.

Villas-Boas, J. M. (2004), 'Price Cycles in Markets with Customer Recognition', The RAND Journal of Economics 35(3), 486–501.

Online Appendix for

Revealed Privacy Preferences: Are Privacy Choices Rational?

Appendix A: Figures

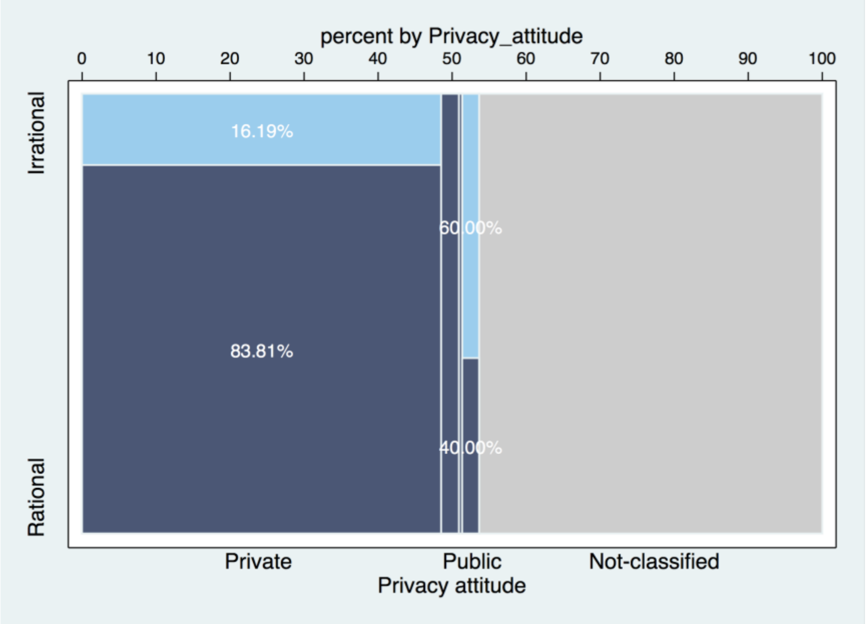

Figure A.1: Proportion of different privacy attitudes

Notes: The privacy-attitude classification used in this figure requires subjects classified as having a specific privacy attitude to choose within 10 percent of the unique point prediction for that attitude for both reports in both two-dimensional problems (within 1 percent of the whole choice set). This strict criterion results in 100 subjects remaining unclassified.

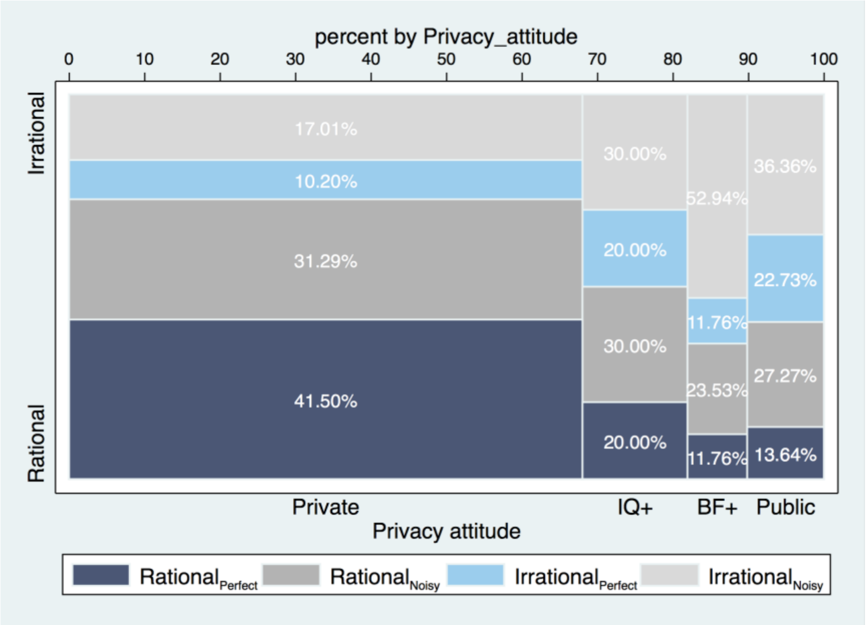

Figure A.2: Rational-Perfect Percentage

Grey regions: subjects pricing privacy control with inconsistency (noisy)

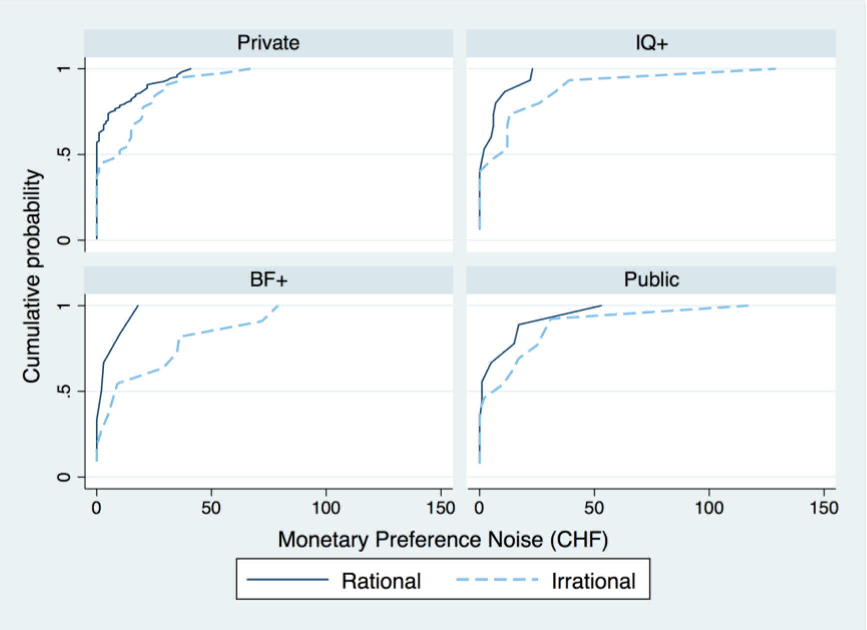

Figure A.3: Cumulative distribution functions of Monetary-value noise for irrational and rational types, by privacy attitudes.

Notes: Irrational types squander more money than rational types, regardless of their privacy attitudes.

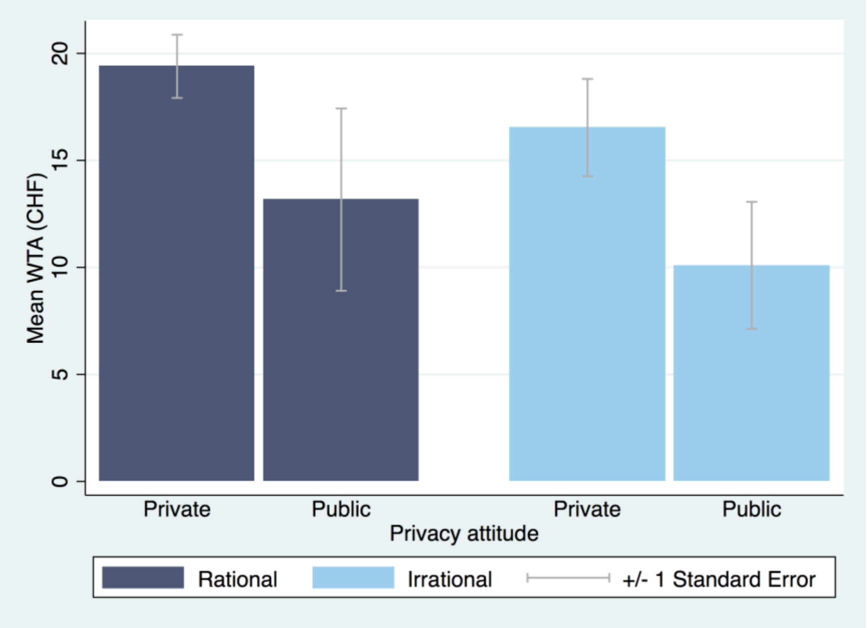

Figure A.4: Willingness-to-accept by Privacy Attitude and Rationality

Notes: Regardless of rationality, private types are willing to pay more for privacy than public types.

Appendix B: Tables

Table A.1: Including Item-dependent Types in Sample When Testing the Robustness of the Correlation between Privacy Attitude and Daily-life Personal Data Sharing (OLS)

	(1)	(2)	(3)	(4)
	Sharing in social networks (Factor 1)			
Public	0.814^{***}	0.788^{***}	0.534^{***}	0.496^{***}
	(3.55)	(3.37)	(3.65)	(3.42)
Controls	No	Yes	No	Yes
Constant	-0.0822	0.0551	-0.178^{*}	-0.0877
	(-1.13)	(0.17)	(-2.11)	(-0.27)
N	198	198	198	198

t statistics in parentheses, * p < .05, ** p < .01, *** p < .001

Notes: Item-dependent type is coded as Public=0 in Model 1 and Model 2; it is coded as Public=1 in Model 3 and Model 4. This table complement the results in Table 5: including the item-dependent type increasing the sample size from 152 to 198.

Table A.2: Correlation between Monetary-value Noise and CCEI

	Monetary-Value Noise			
	Private		Public	
	(1)	(2)	(3)	(4)
e^{*}	17.15^{*}	16.43^{*}	11.96	-9.081
	(2.35)	(2.21)	(0.16)	(-0.11)
Controls	No	Yes	No	Yes
Constant	-9.78	-7.61	3.71	46.63
	(-1.28)	(-0.87)	(0.05)	(0.56)
N	147	147	22	22

Notes: Controls include income(categorical), sex (dummy), education (categorical) and marital status (categorical). t statistics in parentheses, * p < .05, ** p < .01, *** p < .001

Appendix C: Example of Stage 3 Interfaces

Figure A.5 is the screenshot for eliciting the measure "the number of active social media accounts." Notice that the potential payments shown at the bottom of the screen can differ in three cases: "Answers Not Checked" for not having a random check, "Checked & True" for having a random check with *all* answers in Stage 3 verified to be correct, and "Checked & Not True" for having a random check and find any false reporting.

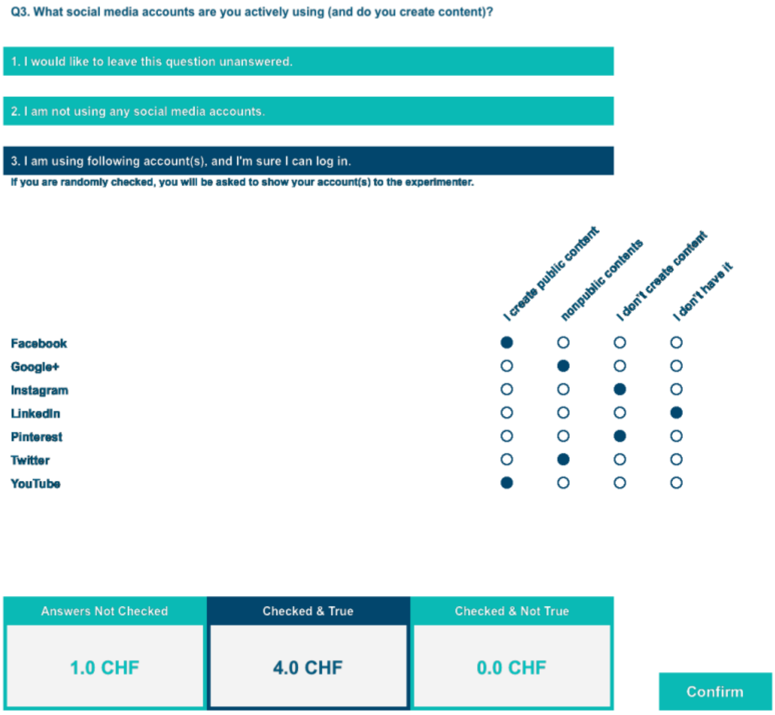

Q3. What social media accounts are you actively using (and do you create content)?

1. I would like to leave this question unanswered.

2. I am not using any social media accounts.

3. I am using following account(s), and I'm sure I can log in.

If you are randomly checked, you will be asked to show your account(s) to the experimenter.

	I create public content	nonpublic contents	I don't create content	I don't have it
Facebook	●	○	○	○
Google+	○	●	○	○
Instagram	○	○	●	○
LinkedIn	○	○	○	●
Pinterest	○	○	●	○
Twitter	○	●	○	○
YouTube	●	○	○	○

Answers Not Checked	Checked & True	Checked & Not True
1.0 CHF	4.0 CHF	0.0 CHF

Figure A.5: Screenshots of the Interface for Stage 3

The decision problem started with a screen that showed only three main answers in colored squares and a blank payment table at the bottom. When choosing the first main answer, "I would like to leave this question unanswered," or the second main answer, "I am not using any social media accounts," a subject earned 1 CHF in all three cases, regardless of whether the random check took place or not, because these two answers are unverifiable.[34] If a subject chose the third answer, "I am using following account(s), and I’m sure I can log in," the screen would then show a table with various social media accounts and choices with different degrees of content sharing behaviors. In this case, Subjects were required to indicate one answer for each social media account and would see an increment of 0.5 CHF of the payoff under the case "Checked and True" when choosing a verifiable answer (either "I create public content," "nonpublic contents" or "I don’t create content" in this example) for a social media account.[35] The "Checked and Not True" payoff was always zero since a subject earned nothing if any answer was found incorrect. Finally, a notice was shown below the main answer to remind subjects that their choices will be verified if the random check takes place. The same procedure was applied to all questions in Stage 3 to incentivize subjects to provide correct information regarding their use of personal information.

[34] Notice that subjects were indifferent between choosing answers 1 and 2 and could only earn additional money when choosing the answer 3. This design gives subjects who used some social media accounts but did not want to be randomly checked a higher incentive to honestly indicate that they would like to leave the question answered, instead of misreporting that they are not using any social media accounts. Allowing subjects to leave a question unanswered leads to data attrition (9 out of 216 subjects in this question, and 18 out of 216 subjects leave at least one question unanswered), but the reported data is expected to be less noisy.

[35] Since the answer “I don’t have it” is not verifiable in the experiment, choosing this answer did not increase the “Checked & True” payment.

Appendix D: Factor Analysis for Daily-life Personal Information Sharing

We conduct the factor analysis using the four variables of real-world data sharing behaviors as inputs, with orthogonal Varimax rotation. The rotated factor loadings, the eigenvalues, the percentages of variance captured by the two factors and the uniqueness of the four variables are summarized in Table A.3.

Following the Kaiser criterion (Kaiser, 1960), we retain these two factors with eigenvalues larger than one that explain a total of 66.59% of the variance. Inspecting the factors provides a clear interpretation. Factor 1 is related to social networking due to its high loadings for *Facebook Public*, the number of public Facebook personal information items (0.8370), and *Social Media Public*, the number of social networking websites for which individuals create public content (0.8271). The social-networking indicator explained 37.89% of the variance. We interpret Factor 2 as representing the exchange of personal information for money or services because of its high loadings for *Memberships*, the number of store club memberships in use (0.8473), and *Location Shared*, sharing one's location with at least some apps (0.6542). The variance explained by Factor 2 is 28.70%.

Table A.3: Factor Analysis for Daily-life Personal Information Sharing Measures

	Rotated Factor Loadings (n=198)		
Variable	Factor1:	Factor2:	Uniqueness
Facebook Public	**0.8370**	-0.0213	0.2990
Social Media Public	**0.8271**	0.0400	0.3143
Memberships	-0.1208	**0.8473**	0.2675
Location Shared	0.3412	**0.6542**	0.4556
Eigenvalue	1.5157	1.1480	
% of Variance	37.89	28.70	
Total Variance			**66.59%**

The uniqueness of our four variables is somewhat low overall, with the variable "*Location Shared*" having a mildly more substantial unique variance (45.56%) that is not shared with other variables. This characteristic indicates that our daily-life personal information sharing measures are indeed related to each other, and this set of variables are legitimate for factor analysis.

Finally, we create two standardized indicators of daily-life information sharing domains for each subject using the regression scoring coefficients in Table A.4.

Table A.4: Scoring Coefficients
(Method: regression; based on varimax rotated factors)

Variable	Factor1	Factor2
Facebook Public	0.5598	-0.0849
Social Media Public	0.5484	-0.0302
Memberships	-0.1475	0.7555
Location Shared	0.1758	0.5491

Appendix E: Collective-level Analysis on Domain-specific Daily-life Information Sharing

In this appendix, we test whether the privacy attitude and the rationality in this study correlate to how people behave in the two domains of daily-life personal data sharing: social-networking (Factor 1) and the exchange of personal information for money or services (Factor 2) at the collective level. We find that people classified as the public type in the lab share more personal information publicly in their daily life than the private type in the social-networking domain but not in the domain of the exchange of money and services. Figure A.6(a) shows that this trend holds for both rational and irrational types. Rank-sum tests confirm that public types share significantly more than private types in the social- networking domain for both rational and irrational subjects at the 5% level (z= −2.170 for rational types and z= −2.312 for irrational types).[36] On average, pooling across rational and irrational types, the public type scores 0.7317 in the social-networking domain. A t-test adjusted for the unequal variances shows that this score is significantly higher than the mean score of the private types: −0.1781 at 1% level (t = −3.2347). A rank-sum test further shows the difference between the scores of these two types is significant at 0.1% level (z = −3.285). Hence, public types share significantly more than private types in the social-networking domain.

However, Figure A.6(b) indicates that public and private types exhibit no systematic significant differences when sharing personal data for the exchange of money and services. Nevertheless, we observe some evidence that people who are irrational when deciding the privacy levels for different information items in Part 1 give over more personal information for money or services than the rational types. Figure A.6(b) shows this trend for both the private and public types. On average, the mean score of sharing personal information in exchanging for money or services for the *irrational* types (pooling private, IQ+, BF+ and public) is 0.1607. However, this score is only marginally significantly higher than the score for the *rational* types: −0.0918 at the 10% level (t = −1.6712, two-tail),[37] so we hesitate to make too much of this relationship.

[36] Due to the nature of the limited number of observations of public types (with 9 observations of rational-public type and 11 observations of irrational-public type), we employ rank-sum tests to compare the value difference between public and private types.

[37] A rank-sum test confirms a similar result (z = −1.606).

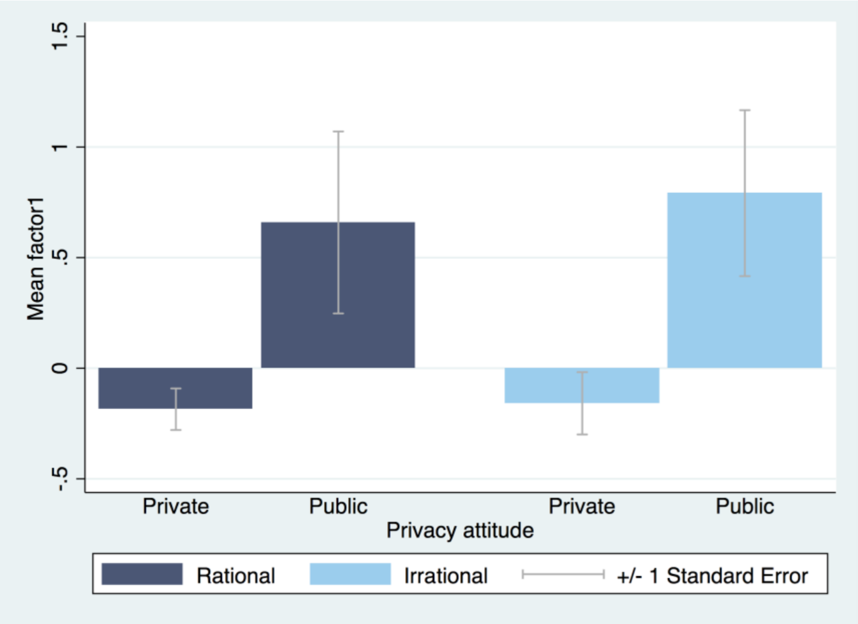

(a) Social networking

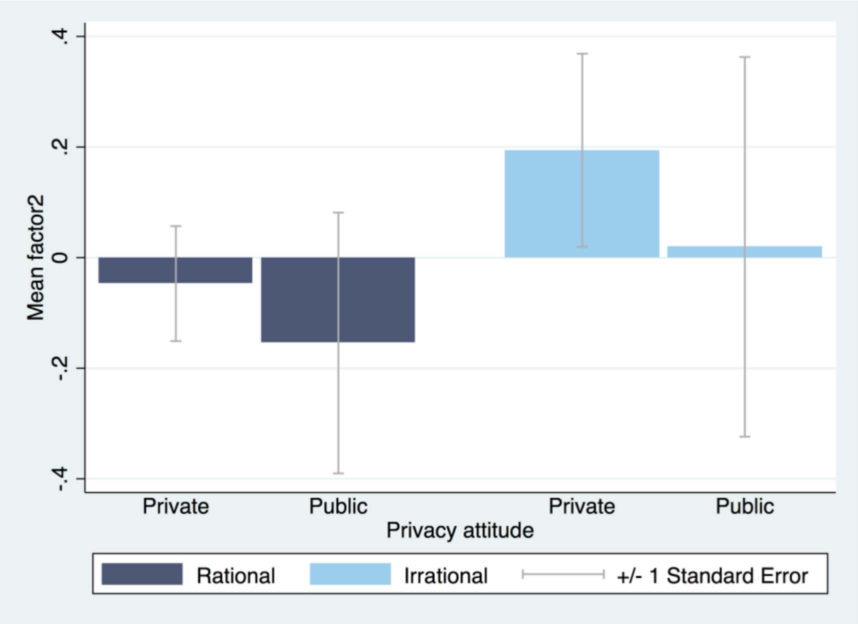

(b) Exchange for money or services

Figure A.6: Domain-specific daily life information sharing

Appendix F: Collective-level Analysis on General Privacy Attitude

When comparing the general privacy attitude between subjects with different choice-based privacy attitudes, Figure A.7 shows that, on average, the private type of both the rational and irrational categories considers general privacy more important than the corresponding public types. Overall, a rank-sum test shows that this difference in self-reported privacy attitudes between the private and the public types is strongly significant at the 0.01% level (z = 4.178). This difference is significant for both the rational types (significant at 1 % level, with z = 3.215) and for the irrational types (significant at 5 % level, with z = 2.484).

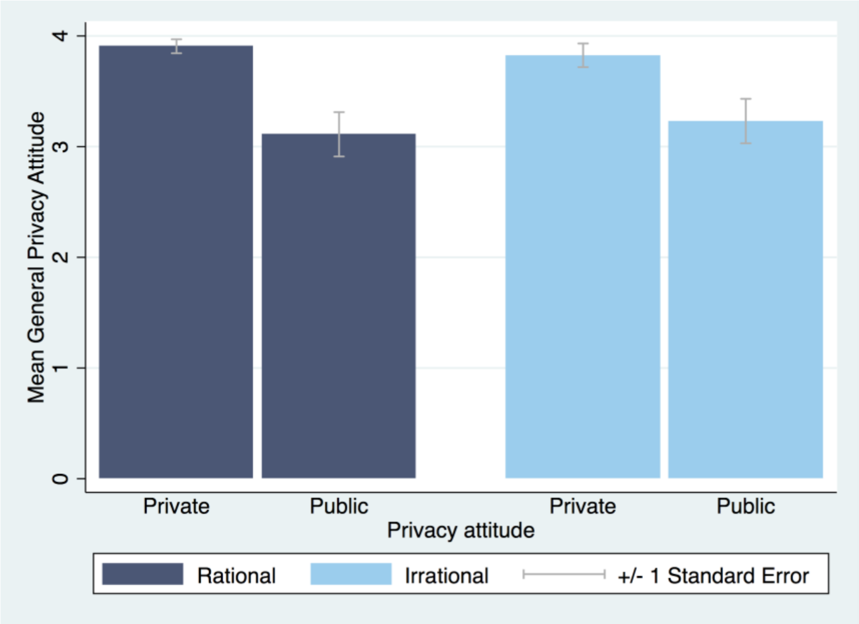

Figure A.7: Self-reported General Privacy Attitudes

Appendix G: Rationality when Making Monetary Choices

Appendix G.1 Classification of Monetary Attitudes

We classify each of our subjects in the follow-up condition into one of four types of monetary attitudes: valuing money (preferring less monetary loss in both accounts), not-valuing-money (preferring more monetary loss in both accounts, (D)) and two account-dependent types (only valuing money in the BLUE account (B) and only valuing money in the RED account (R)). Using subjects' choices in the unrestricted two-dimensional money-attitude questions, we classify each of our subjects into the monetary attitude that generates predictions that are closest to the subject's choices in the two-dimensional monetary attitude questions.

Unlike privacy attitudes, people's monetary attitudes are much more homogenous. The majority (91 out of 94) of subjects in the money condition is classified as valuing money. Figure A.8(a) shows the percentage of subjects of valuing-money types and other types money attitudes, classified as described above, and the proportions of each type who exhibit consistency with GARP. When making monetary choices, 97 percent of subjects value money in both accounts. The remainder are presented as "Others," consisting of around 3 percent of the R type and 1 percent of the not-valuing-money type. Thus, most of our subjects prefer less monetary loss in both accounts in this context, which is consistent with our expectations.

Appendix G.2 Conformity with GARP

We adapt the rationality test according to the different monetary attitudes—i.e., the direction of preferences on the monetary losses in accounts. We find in total, around 44 percent of subjects exhibit perfect consistency with GARP. As shown in Figure 3(a), the degree of rationality is the highest among individuals we separately classify as valuing-money types: more than 45 percent of the valuing-money subjects are rational. The rational percentages corresponding to other types (R and the not-valuing-money types) are both 0 percent.

Figure A.8(b) shows the e* scores, arranged in increasing order, for individuals with different monetary attitudes. The average e* scores for each type are also displayed in Table 2. This score averaged 1.0599 for the valuing-money types (97% of subjects in this condition), which implies that, on average, the monetary budget needs to be reduced by about 6 percent to eliminate all GARP violations by a valuing-money subject. This score is not statistically different from that in our privacy condition.

Among valuing-money types, the proportion of rational subjects in this follow-up study (ranging from 45.1 to 78.0 percent, depending on whether one uses the 0.95 CCEI threshold proposed by Varian (1991) to allow for minor violations)[38] is more similar to that in most of the previous studies included in Table 2, compared to our privacy study. Given this follow-up condition is studying monetary tradeoffs that were studied in earlier studies, it is not surprising we see this pattern. Similar to the results in our privacy condition, the proportion of individuals who allocate privacy levels rationally in our study does not seem to differ substantially from the rational proportion of individuals facing economic tradeoffs in earlier studies.

[38] Depending on whether we require full rationality (zero GARP violations) or allow small violations that forfeit less than 5 percent of their privacy/monetary budgets (the 0.95 threshold of CCEI proposed in Varian (1991)).

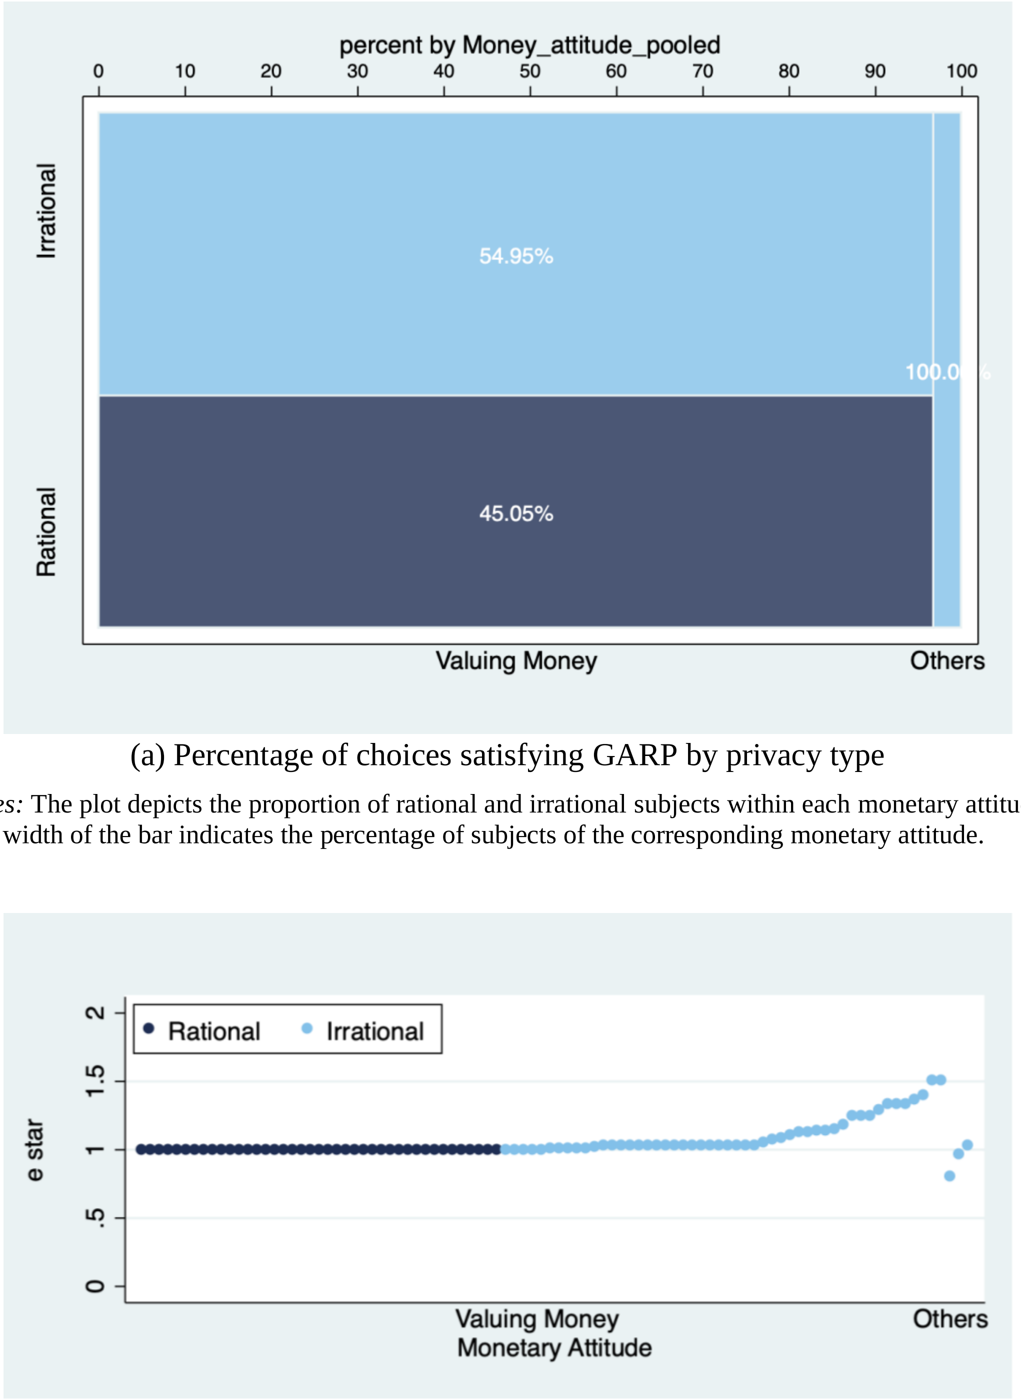

(a) Percentage of choices satisfying GARP by privacy type

Notes: The plot depicts the proportion of rational and irrational subjects within each monetary attitude. The width of the bar indicates the percentage of subjects of the corresponding monetary attitude.

(b) Efficiency score: CCEI (e^*)

Notes: The scatterplot depicts the distribution of the efficiency measure, e^*, for subjects with different monetary attitudes. Each dot represents the e^* of one subject. Within each attitude, rational subjects (subjects with no GARP violations) are plotted first with navy dots. The e^* of irrational subjects are sorted in ascending order, represented by light blue dots.

Figure A.8: Rationality Measures for Monetary Choices

Appendix H: Instructions

Appendix H.1 Privacy Condition

Welcome

This experiment is about decision making. You will be paid for your participation, and the amount of money you earn will be determined by your decisions and also partly by chance. The entire experiment lasts no more than 2 hours.

For completing this experiment, you receive a standard payment of CHF 40. However, you might add to this payment during this experiment. Your final payment will range from CHF 40 to possibly more than CHF 100. Please pay attention to the instructions; this will help you make sound decisions and earn a greater amount of money.

The experiment has three stages. You will now proceed to Stage 1. As in every experiment, you may stop participating at any point. However, if you decide to suspend your participation, you will only receive CHF 20 after completing the shorter version of this experiment.

Stage 1: Specifying 20 most preferred uses of your personal information

In this stage, you will see 20 decision problems. In each decision problem you will specify your most preferred use of the personal information contained in your assessment reports, in the form of "report-viewing combinations."

In each decision problem, you will see different possible numbers of viewers of your reports, by using your slider or sliders on the screen. Your task is to choose your most preferred use of your personal information from all the possibilities. Each of the combinations chosen in Stage 1 will be equally likely to be selected to determine the use of your personal information at the end of the experiment, regardless of what you choose. Moreover, the numbers of viewers you choose will NOT influence the randomly drawn offer price for selling your personal information. Therefore, it is in your best interest to select the report-viewing combination that you would most prefer to have implemented, rather than trying to select a different one in the hope that it will give you more money.

Specifically, choosing your most preferred combination from all the possibilities will give you the following two advantages. First, you can guarantee that your most preferred combination will be one of the 40 combinations that may determine the use of your personal information at the end of the experiment. Second, the more you prefer a combination, the less likely you will be willing to give up money to avoid having your information used in that manner. Hence, you can expect to earn more money at the end of the experiment by selecting the combination that you most prefer in every case.

This stage consists of two parts: Stage 1A and Stage 1B. The choice interface will work slightly differently in each part. Please pay attention to the detailed instructions on how the interface works in each part.

In the following screens, you will be given a quiz to make sure you understand how Stage 1 works before starting it. Please answer each of the questions before proceeding.

Stage 1A: Specifying preferred uses of your personal information

This part consists of 4 decision problems. In each decision problem, you choose your most preferred use of your personal information, separately for EACH report. That is, you choose the number of viewers that you most prefer to view each assessment report, independently of how many will view the other report. Once you set the number of viewers for each report, this will create a report-viewing combination that corresponds to a possible use of your personal information.

Two examples of your decision screen in this part appear below.

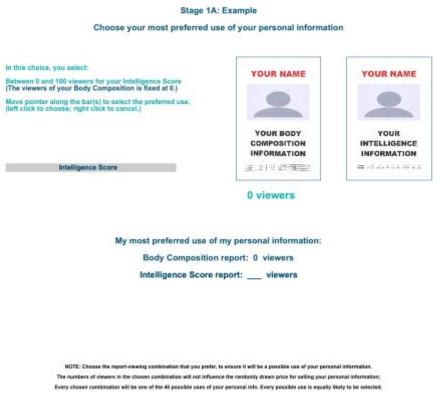

Stage 1A: Example

Choose your most preferred use of your personal information

In this choice, you select:

Between 0 and 100 viewers for your Intelligence Score
(The viewers of your Body Composition is fixed at 0.)

Move pointer along the bar(s) to select the preferred use.
(left click to choose; right click to cancel.)

Intelligence Score
YOUR NAME
YOUR BODY COMPOSITION INFORMATION
YOUR NAME
YOUR INTELLIGENCE INFORMATION
0 viewers

My most preferred use of my personal information:

Body Composition report: 0 viewers

Intelligence Score report: ___ viewers

NOTE: Choose the report-viewing combination that you prefer, to ensure it will be a possible use of your personal information.
The numbers of viewers in the chosen combination will not influence the randomly drawn price for selling your personal information;
Every chosen combination will be one of the 40 possible uses of your personal info. Every possible use is equally likely to be selected.

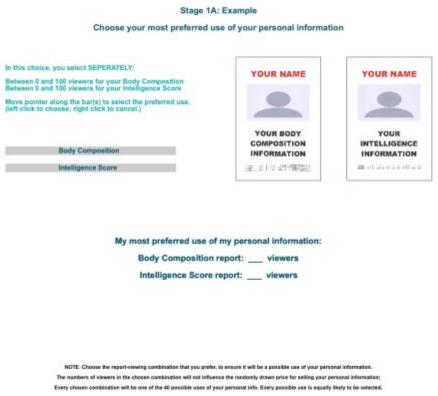

Stage 1A: Example

Choose your most preferred use of your personal information

In this choice, you select SEPERATELY:

Between 0 and 100 viewers for your Body Composition
Between 0 and 100 viewers for your Intelligence Score

Move pointer along the bar(s) to select the preferred use.
(left click to choose; right click to cancel.)

Body Composition
Intelligence Score
YOUR NAME
YOUR BODY COMPOSITION INFORMATION
YOUR NAME
YOUR INTELLIGENCE INFORMATION

My most preferred use of my personal information:

Body Composition report: ___ viewers

Intelligence Score report: ___ viewers

NOTE: Choose the report-viewing combination that you prefer, to ensure it will be a possible use of your personal information.
The numbers of viewers in the chosen combination will not influence the randomly drawn price for selling your personal information;
Every chosen combination will be one of the 40 possible uses of your personal info. Every possible use is equally likely to be selected.

- The example on the left corresponds to the case in which the number of viewers of your Body Composition report is fixed at 0 viewers, and you can choose between 0 and

100 viewers for your Intelligence Score report. This means that your Body Composition report will not be shown to any viewers if this report-viewing combination is selected to determine the use of your personal information. In this case, you do not see the bar corresponding to your Body Composition report and you would only choose your preferred number of viewers for your Intelligence Score report.

- The example on the right corresponds to the case in which you can choose SEPARATELY between 0 and 100 viewers for your Body Composition report, and between 0 and 100 viewers for your Intelligence Score report.

To start a decision problem, move your mouse to any point on one of the slider bars. Then the number of viewers for the indicated report, corresponding to that point, will appear. Move your pointer along each bar to explore available choices and find your preferred number of viewers for the corresponding report. Notice that the numbers of bars, the numbers of viewers on the bars, and how the numbers of viewers for each report increase or decrease from point to point may vary across decision problems.

When you are ready to make your decision, please left click your mouse on the corresponding point on the slider bar to indicate your most preferred number of viewers for that report. Then, your preferred report-viewing combination will be automatically entered in the answer space on your screen. Below, you see two examples.

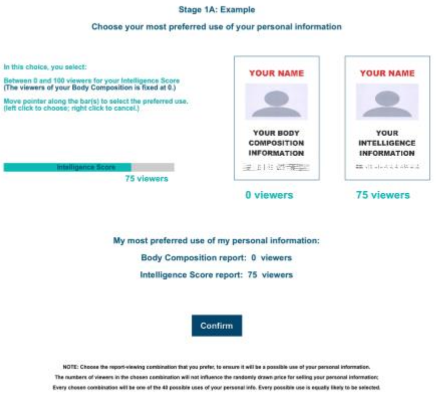

Stage 1A: Example

Choose your most preferred use of your personal information

In this choice, you select:

Between 0 and 100 viewers for your Intelligence Score
(The viewers of your Body Composition is fixed at 0.)

Move pointer along the bar(s) to select the preferred use.
(left click to choose; right click to cancel.)

My most preferred use of my personal information:

Body Composition report: 0 viewers

Intelligence Score report: 75 viewers

NOTE: Choose the report-viewing combination that you prefer, to ensure it will be a possible use of your personal information.
The numbers of viewers in the chosen combination will not influence the randomly drawn price for selling your personal information;
Every chosen combination will be one of the 40 possible uses of your personal info. Every possible use is equally likely to be selected.

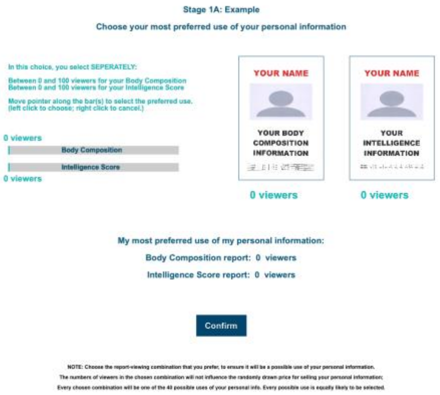

Stage 1A: Example

Choose your most preferred use of your personal information

In this choice, you select SEPERATELY:

Between 0 and 100 viewers for your Body Composition
Between 0 and 100 viewers for your Intelligence Score

Move pointer along the bar(s) to select the preferred use.
(left click to choose; right click to cancel.)

My most preferred use of my personal information:

Body Composition report: 0 viewers

Intelligence Score report: 0 viewers

NOTE: Choose the report-viewing combination that you prefer, to ensure it will be a possible use of your personal information.
The numbers of viewers in the chosen combination will not influence the randomly drawn price for selling your personal information;
Every chosen combination will be one of the 40 possible uses of your personal info. Every possible use is equally likely to be selected.

- The example on the left corresponds to the case in which the number of viewers of your Body Composition report is fixed at 0 and you choose 75 viewers for your Intelligence Score report. This decision means that you prefer having 75 viewers for your Intelligence Score report over ANY other available numbers of viewers in this decision problem. This includes ALL numbers of viewers greater than 75 (such as 100) and ALL numbers of viewers fewer than 75 (such as 0).

- The example on the right corresponds to the case in which you choose 0 viewers for your Body Composition report and 0 viewers for your Intelligence Score report. This decision means that you prefer having 0 viewers for your Body Composition report and Intelligence Score report over ANY other available numbers of viewers in this decision problem. This includes ALL numbers of viewers greater than 0 (such as 100), for either report.

Of course, the numbers of viewers that you select separately for each report depends on how you want your personal information to be used. You are free to make any choice in the range of possible choices separately for each report. However, as we note earlier, it is in your best interest to always select the combination that you find most preferable in all decision problems.

If you want to cancel your choice, please right click your mouse to continue exploring available choices.

After choosing your most preferred numbers of viewers in a decision problem, confirm your decision by clicking on the "confirm" button. You will then proceed to the next decision problem.

For each of the first two decision problems in Stage 1A, you will see the decision screen twice. We do this to make sure everyone understands the choices they are making, before proceeding to the remaining decision problems. In the first two decision problems, you will see the decision screen and make your choice. After you select a combination, an extra screen will then explain to you, in detail, the possible consequences of your choice. You will then return to the same decision screen in order to enter your final choice. This will be the procedure for the first and second decision problems, regardless of what you choose. However, for the remainder of the decision problems in Stage 1, you will only encounter each decision problem once and cannot change your answer after leaving the decision screen. So please choose carefully and finalize your decision before clicking the confirm button.

Note: The numbers of viewers you choose for a combination will not influence the randomly drawn offer price for selling the use of your personal information.\n Therefore, it is in your best interest to select the report-viewing combination that you would most prefer to have implemented,\n rather than trying to select a different one in the hope that it will give you more money.

Stage 1B: Specifying preferred uses of your personal information

This part consists of 16 decision problems. In each decision problem, you now choose your most preferred use of your personal information, jointly for BOTH reports. That is, you simultaneously choose the combination of numbers of viewers that you most prefer to view each assessment report. Once you set the numbers of viewers for both reports, this will create a report-viewing combination that corresponds to a possible use of your personal information.

An example of your decision screen in this part appears below.

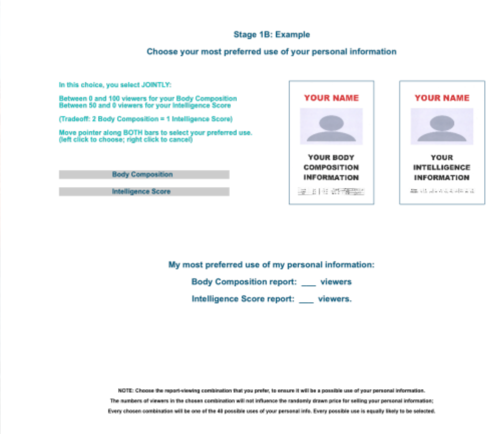

Stage 1B: Example

Choose your most preferred use of your personal information

In this choice, you select JOINTLY:

Between 0 and 100 viewers for your Body Composition
Between 50 and 0 viewers for your Intelligence Score

(Tradeoff: 2 Body Composition = 1 Intelligence Score)

Move pointer along BOTH bars to select your preferred use.
(left click to choose; right click to cancel)

Body Composition
Intelligence Score
YOUR NAME
YOUR BODY COMPOSITION INFORMATION
YOUR NAME
YOUR INTELLIGENCE INFORMATION

My most preferred use of my personal information:

Body Composition report: ___ viewers

Intelligence Score report: ___ viewers.

NOTE: Choose the report-viewing combination that you prefer, to ensure it will be a possible use of your personal information.
The numbers of viewers in the chosen combination will not influence the randomly drawn price for selling your personal information;
Every chosen combination will be one of the 48 possible uses of your personal info. Every possible use is equally likely to be selected.

In this example, you can choose JOINTLY between 0 and 100 viewers for your Body Composition report, and between 50 and 0 viewers for your Intelligence Score report. Notice that, in this example, reducing the number of viewers for your Body Composition report by 2 increases the number of viewers for your Intelligence Score report by 1. This tradeoff rate may vary across decision problems, but your decision screen will always inform you of the tradeoff for every problem.

To start a decision problem, move your mouse to any point on one of the slider bars. Then the numbers of viewers for BOTH reports, corresponding to that point, will appear. Move your pointer along bars to explore available choices and find your most preferred combined numbers of viewers for the corresponding reports. Remember that the numbers of viewers on the bars and how the numbers of viewers for each report increase or decrease from point to point may vary across decision problems.

When you are ready to make your decision, please left click your mouse on the corresponding point on the slider bars to indicate your most preferred combined numbers of viewers for both reports. Then, your preferred report-viewing combination will be automatically entered in the answer space on your screen.

Below, you see an example that corresponds to the case in which you choose 36 viewers for your Body Composition report and 32 viewers for your Intelligence Score report.

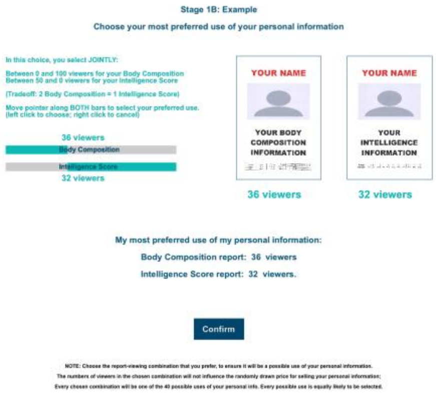

Stage 1B: Example

Choose your most preferred use of your personal information

In this choice, you select JOINTLY:

Between 0 and 100 viewers for your Body Composition
Between 50 and 0 viewers for your Intelligence Score

(Tradeoff: 2 Body Composition = 1 Intelligence Score)

Move pointer along BOTH bars to select your preferred use.
(left click to choose; right click to cancel)

My most preferred use of my personal information:

Body Composition report: 36 viewers

Intelligence Score report: 32 viewers.

NOTE: Choose the report-viewing combination that you prefer, to ensure it will be a possible use of your personal information.
The numbers of viewers in the chosen combination will not influence the randomly drawn price for selling your personal information;
Every chosen combination will be one of the 40 possible uses of your personal info. Every possible use is equally likely to be selected.

This decision means that you prefer having the combination with 36 viewers for your Body Composition report and 32 viewers for your Intelligence Score report over ANY other available combination in this decision problem. This includes ALL combinations that consist of more viewers for your Body Composition report and fewer viewers for your Intelligence Score report (such as 100 viewers for your Body Composition report and 0 viewers for your Intelligence Score report) and ALL combinations that consist of fewer viewers for your Body Composition report and more viewers for your Intelligence Score report (such as 0 viewers for your Body Composition report and 50 viewers for your Intelligence Score report).

Of course, the numbers of viewers that you select jointly for both reports depends on how you want your personal information to be used. You are free to make any choice in the range of possible combined choices for both reports. However, as we note earlier, it is in your best interest to always select the combination that you find most preferable in all decision problems.

If you want to cancel your choice, please right click your mouse to continue exploring available choices.

After choosing your most preferred numbers of combined viewers in a decision problem, confirm your decision by clicking on the "confirm" button. You will then proceed to the next decision problem.

Remember that for all of the remaining 16 decision problems, you will only see the decision screen once. You cannot change your answer after leaving the decision screen. So please choose carefully and finalize your decision before clicking the confirm button.

Note: The numbers of viewers you choose for a combination will not influence the randomly drawn offer price for selling the use of your personal information.\n Therefore, it is in your best interest to select the report-viewing combination that you would most prefer to have implemented,\n rather than trying to select a different one in the hope that it will give you more money.

Stage 2: Reviewing and selling the use of your personal information

In this stage, you will review all 40 report-viewing combinations that may be selected to determine the use of your personal information contained in the assessment reports. You chose 20 combinations in Stage 1, and the computer chose the other 20. Each of the 40 report-viewing combinations you will encounter in Stage 2 is equally likely to be selected to determine the use of your personal information at the end of this experiment.

In each decision problem, you will specify which prices you are willing and not willing to accept in exchange for selling your personal information so that it is used in the manner specified by the report-viewing combination. Specifically, you will indicate the lowest price that you are willing to accept, meaning that you accept that price or any price that is higher.

In each decision problem, you will see a possible use of your personal information: how many random viewers at UZH or ETH will see each of your reports. You will then specify the lowest price, between CHF 0 and CHF 40 that you are willing to accept to sell your personal information and have it used in this manner. If you are unwilling to accept any price between CHF 0 and CHF 40, you may also indicate that the lowest price is "more than CHF 40," which guarantees that you will not sell your personal information at any possible price.

To choose a specific amount, move your pointer to the corresponding point on the slider bar. This will highlight all of the values that are at least as high as the value you select, which means you consider any price offer in this range to be acceptable. When you are ready to make your decision, left click on the corresponding point on the bar, and the amount will be automatically entered in the answer space on your screen. If you want to cancel your choice, please right click your mouse.

The use of your personal information

At the end of this experiment, one of the 40 report-viewing combinations in Stage 2 will be randomly selected. Every combination in Stage 2 is equally likely to be selected. The computer will also randomly select a price to offer you for using your personal information according to that combination. All prices between CHF 0 and CHF 40 are equally likely to be selected. The lowest acceptable price that you indicated for the selected report-viewing combination and the randomly drawn offer price will determine whether or not you sell your personal information and allow it to be used in the manner specified in that combination.

- If the randomly drawn offer price is lower than the lowest acceptable price you specified,

then you will not sell your personal information to be used as specified in the selected report-viewing combination. The randomly drawn offer price will not be added to your payment for this experiment, and you will shred all three envelopes with your assessment reports and pictures in them at the end of the experiment.

- If the randomly drawn offer price is greater than or equal to the lowest acceptable price you specified,

then you will sell your personal information to be used as specified in the selected report-viewing combination. The randomly drawn offer price will be added to your payment for this experiment, and your assessment report(s), with your name, signature and profile photo will be shown to the corresponding number of UZH or ETH viewers.

NOTE

In the following screens for Stage 2, you will specify the lowest price at which you are willing to sell your personal information contained in the assessment reports and allow it to be used as specified in the presented report-viewing combination. Please choose the lowest acceptable prices carefully, since one of them will determine whether or not your personal information will be sold and used in the manner specified by the report-viewing combination and whether the randomly drawn offer price will be added to your payment.

Also note that your decision of the lowest price you are willing to accept does not affect the randomly drawn offer price that you end up receiving, only whether or not you end up being paid the randomly drawn offer price in exchange for the use of your personal information. Therefore, it is in your best interest to honestly indicate the smallest amount of money, such that you would like to sell your personal information when receiving a price greater than or equal to that amount and would like to receive zero and not sell your personal information for any lower price. Misrepresenting the lowest price you are willing to accept may lead to outcomes that you regret.

For example, suppose your actual lowest acceptable price for a certain use is CHF 10, but you instead state that the lowest price you are willing to accept is CHF 30. Then, if the computer randomly draws a price of CHF 20, you will end up not receiving this amount and not selling your personal information, even though you would have preferred to sell it for any price greater than CHF 10, including a price of CHF 20.

As another example, suppose your actual lowest acceptable price for a certain use is CHF 30, but you instead state that the lowest price you are willing to accept is CHF 10. Then, if the computer randomly draws a price of CHF 20, you will end up receiving CHF 20 and selling your personal information, even though you would have preferred not to sell it for any price below CHF 30, including a price of CHF 20.

In the following screens, you will be given a quiz to make sure you understand how Stage 2 works before starting it. Please answer each of the questions before proceeding.

Options for selling the use of your personal information

In Stage 2, you will review all 40 report-viewing combinations that may be selected to determine the use of your personal information contained in the assessment reports. You can indicate the lowest price you are willing to accept for selling your personal information to be used as specified in each of the possible report-viewing combinations.

You may decide to skip this stage. If you choose to skip Stage 2, you give up the opportunity to review possible uses of your personal information and specify the lowest acceptable price for selling your personal information. Hence, the computer will not present report-viewing combinations to you and will specify CHF 0 as the minimum amount for any potential report-viewing combinations in Stage 2. That is, you essentially give permission for any of the 40 possible uses of your personal information that the computer may select from Stage 2, in return for receiving any potential randomly drawn offer price between CHF 0 and CHF 40.

Please indicate which option for determining the use of your personal information you would like to have. The rest of this experiment will proceed accordingly.

□ I would like to complete Stage 2. I care about having control over selling the use of my personal information.

I hereby only give permission to sell the use of my personal information for which I will receive at least the minimum amount I specify in Stage 2, after reviewing how my information will be used in every possible case.

□ I would like to skip Stage 2. I do not care about having control over selling the use of my personal information.

I hereby give permission to sell the use of my personal information for which I will receive an amount greater than or equal to CHF 0, without knowing in advance how my information will be used in all possible cases.

Stage 3: Earning Bonus Payments by Answering Verifiable Questions

In this final stage, you have the opportunity to earn bonus payments by answering verifiable questions regarding your real-life behavior. The bonus payments you receive will depend on your answers and on whether or not the experimenters can verify the truthfulness of your responses. At the end of the experiment, you will be asked to toss a 10-sided die. If you roll a zero, then a random check will be conducted to verify all your answers. Please answer all questions honestly, since you will earn the highest bonus payments when doing so.

For each question, the bottom of the screen will display your bonus payments if you are not checked (if you roll 1 through 9) and if you are checked for honesty (if you roll 0). You earn CHF 1 for answering each main question honestly. For each additional component in your answer that can be verified to be true at the end of the experiment, you earn an additional CHF 0.5 if you are randomly checked and that particular component is verified to be true. However, if you are caught providing any dishonest answer for a question (or any answer that cannot be verified), then you forfeit all bonus payments from this stage. Hence, to earn the highest possible bonus payment, please provide as many verifiable components as possible and answer questions honestly.

Questionnaire

You are now close to the end of this experiment. Please complete a brief questionnaire before receiving your payment.

For questions without a confirm button, please type your answer and it will appear on the screen. Press the "Enter" key for a line break and press the "Esc" key to finish answering and move on to the next question.

Appendix H.2 Abbreviated Experiment

The instructions for the abbreviated experiment consist of the welcome screen below and instructions for Stages 3 and 4 in Appendix H.1

Welcome

This experiment is about decision making. You will be paid for your participation, and the amount of money you earn will be determined by your decisions and also partly by chance. The entire experiment should last no more than 1 hour.

For completing this experiment, you receive a standard payment of CHF 20. However, you might add to this payment during this experiment. Your final payment will range from CHF 20 to CHF 20+bonus payments. Please pay attention to the instructions; this will help you make sound decisions and earn a greater amount of money.

Appendix H.3 Monetary Condition

Welcome

This experiment is about decision making. You will be paid for your participation, and the amount of money you earn will be determined by your decisions and also partly by chance. The entire experiment lasts no more than 1 hour.

For completing this experiment, you receive a standard initial payment of CHF 40. During the experiment, the majority of your earnings will be denominated in points, with an exchange rate of CHF 1 = 5 points.

The points you receive in this experiment will be placed in two accounts: a BLUE account and a RED account. Both accounts have a starting balance of 200 points in them. At the end of the study, the computer will randomly select only one of the two accounts to count for your payment. Each account is equally likely to be selected, but you will not know which account will be the one that is selected until the end of the study. During the experiment, you will make choices that may subtract points from your two accounts. You will also make some choices that might give you additional bonus earnings.

Your final payment will range from CHF 20 to possibly CHF 60. Please pay attention to the instructions; this will help you make sound decisions and earn a greater amount of money.

The experiment has two stages. You will now proceed to Stage 1.

Stage 1: Specifying 20 most preferred distributions of losses

In this stage, you will see 20 decision problems. In each decision problem, you will distribute losses between your two accounts: the BLUE account and the RED account. When facing a decision problem, you can see all possible distributions of possible losses for that decision problem in points by using your slider or sliders on the screen. Your task is to choose your most preferred distribution of possible losses from all the possibilities.

Your payments from this stage will be determined as follows. At the end of the experiment, the computer will randomly select one decision problem and will randomly select one of the two accounts, the BLUE account or the RED account. Each decision problem in Stage 1 will be equally likely to be selected, regardless of what you choose.

Within the selected decision problem, both accounts have an equal probability of being selected. According to your choice in the selected decision problem, the points lost that you distribute to the account that is selected will determine your payment for the experiment. The points lost that you distribute to the other account will not influence your earnings.

This stage consists of two parts: Stage 1A and Stage 1B. The choice interface will work slightly differently in each part. Please pay attention to the detailed instructions on how the interface works in each part.

Stage 1A: Specify possible losses separately for each account

This part consists of 4 decision problems. In each decision problem, you choose your most preferred distribution of possible losses, separately for EACH account. That is, you choose the possible loss in points for your BLUE account independently of the possible loss in points for your RED account.

Two examples of your decision screen in this part appear below.

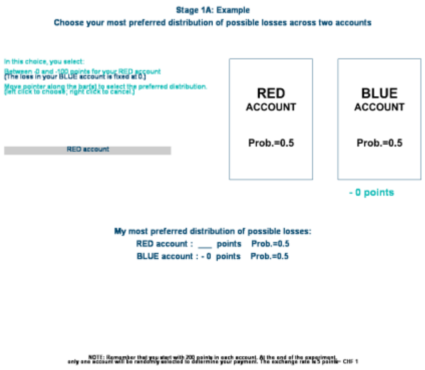

Stage 1A: Example
Choose your most preferred distribution of possible losses across two accounts

In this choice, you select:
Between -0 and -100 points for your RED account
(The loss in your BLUE account is fixed at 0.)
Move pointer along the bar(s) to select the preferred distribution.
(left click to choose; right click to cancel.)

My most preferred distribution of possible losses:
RED account : ___ points Prob.=0.5
BLUE account : - 0 points Prob.=0.5

NOTE: Remember that you start with 200 points in each account. At the end of the experiment, only one account will be randomly selected to determine your payment. The exchange rate is 5 points= CHF 1

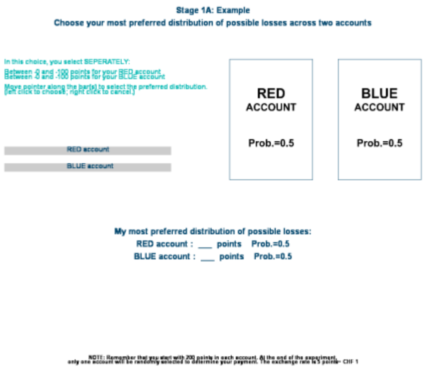

Stage 1A: Example
Choose your most preferred distribution of possible losses across two accounts

In this choice, you select SEPERATELY:
Between -0 and -100 points for your RED account
Between -0 and -100 points for your BLUE account
Move pointer along the bar(s) to select the preferred distribution.
(left click to choose; right click to cancel.)

My most preferred distribution of possible losses:
RED account : ___ points Prob.=0.5
BLUE account : ___ points Prob.=0.5

NOTE: Remember that you start with 200 points in each account. At the end of the experiment, only one account will be randomly selected to determine your payment. The exchange rate is 5 points= CHF 1

- The example on the left corresponds to the case in which the possible loss for your BLUE account is fixed at 0, and you can choose a possible loss between 0 and 100 points for your RED account. This means that if this decision problem is selected and the computer randomly selects the BLUE account then no points will be subtracted from your BLUE account and you will receive 200 points. If, instead, the computer randomly selects the RED account then the loss in points that you select will be subtracted from your RED account. In this decision problem, you do not see the bar corresponding to your BLUE account and you only choose your preferred possible loss for your RED account. Nevertheless, both accounts have an equal probability of being selected as the account that counts for this decision problem.
- The example on the right corresponds to the case in which you can choose SEPARATELY a possible loss between 0 and 100 points for your BLUE account, and a possible loss between 0 and 100 points for your RED account. Again, both accounts have an equal probability of being selected as the account that counts for this decision

problem. In this case, the loss in the account that is randomly selected would be subtracted from your initial 200 points.

To start a decision problem, move your mouse to any point on one of the slider bars. Then the loss for the indicated account, corresponding to that point, will appear. Move your pointer along each bar to explore available choices and find your preferred possible loss for the corresponding account. Notice that the numbers of bars, the distributions of losses, and how many points are lost for each account along the bars may vary across decision problems.

When you are ready to make your decision, please left-click your mouse on the corresponding point on the slider bar to indicate your most preferred possible losses of points for the two accounts. Then, your preferred distribution of possible losses will be automatically entered in the answer space on your screen. Below, you see two examples.

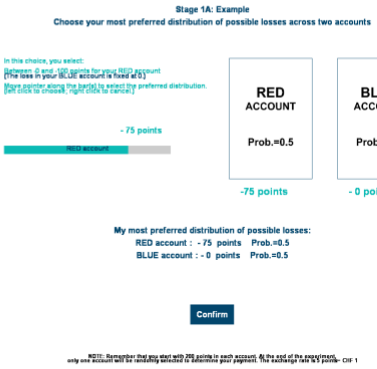

Stage 1A: Example

Choose your most preferred distribution of possible losses across two accounts

In this choice, you select:

Between -0 and -100 points for your RED account

(The loss in your BLUE account is fixed at 0.)

Move pointer along the bar(s) to select the preferred distribution.

(left click to choose; right click to cancel.)

My most preferred distribution of possible losses:

RED account : - 75 points   Prob.=0.5

BLUE account : - 0 points   Prob.=0.5

NOTE: Remember that you start with 200 points in each account. At the end of the experiment, only one account will be randomly selected to determine your payment. The exchange rate is 5 points= CHF 1

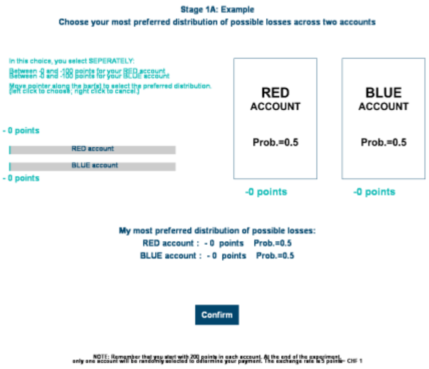

Stage 1A: Example

Choose your most preferred distribution of possible losses across two accounts

In this choice, you select SEPERATELY:

Between -0 and -100 points for your RED account

Between -0 and -100 points for your BLUE account

Move pointer along the bar(s) to select the preferred distribution.

(left click to choose; right click to cancel.)

My most preferred distribution of possible losses:

RED account : - 0 points   Prob.=0.5

BLUE account : - 0 points   Prob.=0.5

NOTE: Remember that you start with 200 points in each account. At the end of the experiment, only one account will be randomly selected to determine your payment. The exchange rate is 5 points= CHF 1

- The example on the left corresponds to the case in which the possible loss for your BLUE account is fixed at 0 points, and you choose a possible loss of 75 points for your RED account. This decision means that you prefer losing 75 points in your RED account over ANY other available choices in this decision problem. This includes ALL losses greater than 75 points (such as losses of 100 points) and ALL losses smaller than 75 points (such as losses of 0 points). If this decision problem is randomly selected, you will receive 200 (200-0) points if the BLUE account is selected to count, or receive 125 (200-75) points if the RED account is selected to count. Both accounts have an equal probability of being selected as the one that counts.
- The example on the right corresponds to the case in which you choose a loss of 0 points for your BLUE account and a loss of 0 points for your RED account. This decision means that you prefer losing 0 points in your BLUE and RED accounts over ANY other available choices in this decision problem. This includes ALL possible

losses greater than 0 (such as losses of 100 points), for either account. If this decision problem is randomly selected, you will receive 200 (200-0) points regardless of which account is selected to count.

If you want to cancel your choice, please right-click your mouse to continue exploring available choices.

After choosing your most preferred distribution of possible losses in a decision problem, confirm your decision by clicking on the "confirm" button. You will then proceed to the next decision problem.

For each of the first two decision problems in Stage 1A, you will see the decision screen twice. We do this to make sure everyone understands the choices they are making, before proceeding to the remaining decision problems. In the first two decision problems, you will see the decision screen and make your choice. After you select a distribution of possible losses, an extra screen will then explain to you, in detail, the possible consequences of your choice. You will then return to the same decision screen in order to enter your final choice. This will be the procedure for the first and second decision problems, regardless of what you choose. However, for the remainder of the decision problems in Stage 1, you will only encounter each decision problem once and cannot change your answer after leaving the decision screen. So please choose carefully and finalize your decision before clicking the confirm button.

In the following screens, you will be given a quiz to make sure you understand how Stage 1 works before starting it. Please answer each of the questions before proceeding.

Stage 1B: Specify possible losses jointly for both accounts

This part consists of 16 decision problems. In each decision problem, you now choose your most preferred distribution of possible losses, jointly for BOTH accounts. That is, you choose the loss in points for your BLUE account and the loss in points for your RED account simultaneously. Remember that if a decision problem is selected to count, the computer will randomly select one of the two accounts as the account that counts for that decision problem, and both accounts have an equal probability of being selected.

An example of your decision screen in this part appears below.

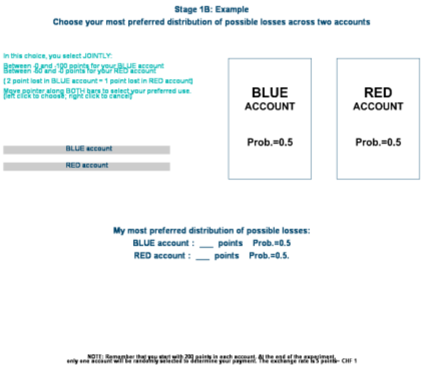

Stage 1B: Example

Choose your most preferred distribution of possible losses across two accounts

In this choice, you select JOINTLY:

Between -0 and -100 points for your BLUE account

Between -50 and -0 points for your RED account

(2 point lost in BLUE account = 1 point lost in RED account)

Move pointer along BOTH bars to select your preferred use.
(left click to choose; right click to cancel)

My most preferred distribution of possible losses:
BLUE account : ___ points Prob.=0.5
RED account : ___ points Prob.=0.5.

NOTE: Remember that you start with 200 points in each account. At the end of the experiment, only one account will be randomly selected to determine your payment. The exchange rate is 5 points= CHF 1

In this example, you can choose JOINTLY a possible loss between 0 and 100 points for your BLUE account, and a possible loss between 50 and 0 points for your RED account. Notice that, in this example, reducing the loss in the BLUE account by 2 points increases the loss in the RED account by 1 point. This tradeoff rate may vary across decision problems, but your decision screen will always inform you of the tradeoff for every problem.

To start a decision problem, move your mouse to any point on one of the slider bars. Then the possible losses for BOTH accounts, corresponding to that point, will appear. Move your pointer along bars to explore available choices and find your preferred distribution of possible losses for the two accounts. Remember that the distributions of losses, and the increase or decrease in points lost in each account as you move the slider may vary across decision problems.

When you are ready to make your decision, please left-click your mouse on the corresponding point on the slider bars to indicate your most preferred distribution of possible losses between the two accounts. Then, your preferred distribution of possible losses will be automatically entered in the answer space on your screen.

Below, you see an example that corresponds to the case in which you choose a possible loss of 36 points for your BLUE account and a possible loss of 32 points for your RED account.

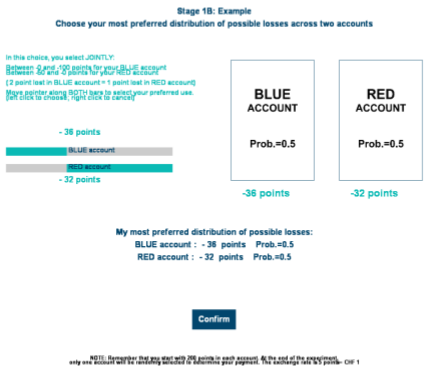

Stage 1B: Example

Choose your most preferred distribution of possible losses across two accounts

In this choice, you select JOINTLY:

Between -0 and -100 points for your BLUE account

Between -50 and -0 points for your RED account

(2 point lost in BLUE account = 1 point lost in RED account)

Move pointer along BOTH bars to select your preferred use.

(left click to choose; right click to cancel)

My most preferred distribution of possible losses:

BLUE account : - 36 points Prob.=0.5

RED account : - 32 points Prob.=0.5

NOTE: Remember that you start with 200 points in each account. At the end of the experiment, only one account will be randomly selected to determine your payment. The exchange rate is 5 points= CHF 1

This decision means that you prefer possibly losing 36 points in your BLUE account and possibly losing 32 points in your RED account over ANY other possible combinations of losses in this decision problem. This includes ALL distributions that consist of a greater possible loss for your BLUE account and a smaller possible loss for your RED account (such as losing 100 points in your BLUE account and losing 0 loss points in your RED account) and ALL distributions that consist of a smaller possible loss for your BLUE account and a greater possible loss for your RED account (such as losing 0 points in your BLUE account and losing 50 in your RED account). If this decision problem is randomly selected, you will receive 164 (200-36) points if the BLUE account is selected to count, or receive 168 (200-32) points if the RED account is selected to count. Both accounts have an equal probability of being selected as the one that counts.

If you want to cancel your choice, please right-click your mouse to continue exploring available choices.

After choosing your most preferred distribution of possible losses in a decision problem, confirm your decision by clicking on the "confirm" button. You will then proceed to the next decision problem.

Remember that for all of the remaining 16 decision problems, you will only see the decision screen once. You cannot change your answer after leaving the decision screen. So please choose carefully and finalize your decision before clicking the confirm button.

Stage 2: Earning Bonus Payments by Answering Verifiable Questions

In the second stage, you have the opportunity to earn bonus payments by answering verifiable questions regarding your real-life behavior. The bonus payments you receive will depend on your answers and on whether or not the experimenters can verify the truthfulness of your responses. At the end of the experiment, you will be asked to toss a 10-sided die. If you roll a zero, then a random check will be conducted to verify all your answers. Please answer all questions honestly, since you will earn the highest bonus payments when doing so.

For each question, the bottom of the screen will display your bonus payments if you are not checked (if you roll 1 through 9) and if you are checked for honesty (if you roll 0). You earn CHF 1 for answering each main question honestly. For each additional component in your answer that can be verified to be true at the end of the experiment, you earn an additional CHF 0.5 if you are randomly checked and that particular component is verified to be true. However, if you are caught providing any dishonest answer for a question (or any answer that cannot be verified), then you forfeit all bonus payments from this stage. Hence, to earn the highest possible bonus payment, please provide as many verifiable components as possible and answer questions honestly.

Questionnaire

You are now close to the end of this experiment. Please complete a brief questionnaire before receiving your payment.

For questions without a confirm button, please type your answer and it will appear on the screen. Press the "Enter" key for a line break and press the "Esc" key to finish answering and move on to the next question.