


ARTICLE

Increasing personal data contributions for the greater public good: a field experiment on an online education platform

Viola Ackfeld¹, Tobias Rohloff² and Sylvi Rzepka^{3*} 

¹University of Cologne, Germany, ²Hasso Plattner Institute, Germany and ³University of Potsdam, Germany

*Correspondence to: E-mail: rzepka@empwifo.uni-potsdam.de

(Received 6 June 2021; revised 4 October 2021; accepted 8 November 2021)

Abstract

Personal data increasingly serve as inputs to public goods. Like other types of contributions to public goods, personal data are likely to be underprovided. We investigate whether classical remedies to underprovision are also applicable to personal data and whether the privacy-sensitive nature of personal data must be additionally accounted for. In a randomized field experiment on a public online education platform, we prompt users to complete their profiles with personal information. Compared to a control message, we find that making public benefits salient increases the number of personal data contributions significantly. This effect is even stronger when additionally emphasizing privacy protection, especially for sensitive information. Our results further suggest that emphasis on both public benefits and privacy protection attracts personal data from a more diverse set of contributors.

Keywords: field experiment; personal data; public good; privacy

Introduction

Personal data have become inputs to algorithms that produce public goods in many different domains. For instance, personal data help predict diseases and their outbreak (Ginsberg *et al.*, 2009; Obermeyer & Emanuel, 2016), which has also triggered the use of tracing apps to mitigate the spread of COVID-19. Personal data can help build algorithms that route traffic more efficiently (Lv *et al.*, 2014; Cramton *et al.*, 2019). Moreover, on open online learning platforms – the setting studied in this article – personal data serve as inputs to algorithms that improve learning experiences (Yu *et al.*, 2017; Fauvel *et al.*, 2018) not only for the contributing user but also for the entire user community. Irrespective of the domain, these examples have something in common: personal data contribute to a public good, be it the ‘absence of disease’ (Fisman & Laupland, 2009), uncongested traffic flow, or free online education. In

© The Author(s), 2021. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

these public good contexts, people share personal data on a voluntary basis to public goods, that is, they ‘donate’ data.¹ Public good providers explicitly use this term themselves. For instance, to create open-source speech-recognition software, Mozilla’s Common Voice project asks to ‘donate your voice.’ Similarly, the online platform openhumans.org asks for donations of personal data to conduct scientific research.^{2,3}

Due to non-rivalry and non-excludability in consumption, contributions of personal data to a public good are as likely to be underprovided as other types of contributions such as money or effort, which have been extensively studied [see, e.g., Frey & Meier (2004), Shang & Croson (2009), Krupka & Croson (2015), or Chaudhuri (2011) for a review]. Yet, tackling the underprovision of personal data to public goods may differ for two reasons: personal data are individual-specific and they are potentially privacy-sensitive. First, individual-specific means that not only the total amount of donations matters but also who and how many different people donate personal data. Imagine algorithms to route a city center’s traffic were missing data on pedestrians (Green, 2019). With these biased input data, we would expect streets to become unsafe. If only a small share of a population truthfully participated in a contact tracing app, then the app may not be very effective in containing the spread of a disease (Nature Editorial, 2020). Likewise in online education, where currently about 80% of MOOC participants fail to reach their intended learning goal (Kizilcec *et al.*, 2020), it is unclear how to best use personal data to build algorithms that support the learning goals of all participants if only 48% of all users share information about themselves – as it is the case at baseline in our setting.⁴ Hence, public good production that uses personal data as inputs to train algorithms requires not only a large but also a diverse and representative database. Otherwise, the quality of the public good may suffer. While other types of individual-specific contributions to public goods have already been studied, for example, knowledge (Zhang & Zhu, 2011; Chen *et al.*, 2020) and feedback (Bolton *et al.*, 2004; Ling *et al.*, 2005; Chen *et al.*, 2010; Cabral & Li, 2015; Bolton *et al.*, 2020), personal data as contributions have so far been neglected. We study whether insights on how to boost other types of contributions transfer to personal data.

Second, personal data contributions may work differently than other individual-specific contributions because compared to providing feedback or knowledge, they come with an additional concern: privacy costs.⁵ These may exacerbate the underprovision of personal data donations to public goods. As Goldfarb and Tucker’s (2012) document, there are certain demographic groups, for example, elderly and women, that are less likely to share personal information. This heterogeneity in privacy

¹Note the difference to contexts in which users ‘pay with their data’ as it would be the case if an app requires personal information provision, making usage excludable.

²See <https://voice.mozilla.org/en> and <https://www.openhumans.org/>.

³Many of these public goods are what Chen *et al.* (2020, p. 3) call ‘digital public good,’ an ‘increasingly important class of public goods,’ but they are not limited to the digital sphere.

⁴Further examples are target services such as planning prompts to learners for whom such services actually can increase certification rate. Planning prompts prove effective only for a subgroup of learners in Andor *et al.* (2018), and the platform can only target this subgroup adequately when having sufficient information to identify this user subgroup.

⁵See Acquisti *et al.* (2016) and Tucker (2015) for comprehensive reviews on the economics of privacy.

concerns challenges the public goods production that takes personal data as inputs. If the elderly are more concerned to share their location, traffic flows cannot be optimized, and contact tracing apps will not be able to serve this at-risk group. Similarly, if users who do not belong to the median user group are less inclined to share information, for example, female users in tech-related online learning, these platforms are more likely to continue to cater to optimize the learning experience for the male median user.

In this article, we study how to increase both the quantity and quality, that is, the diversity of personal data contributions to a public good in light of the public benefits and private privacy costs of sharing such information. More precisely, we conduct a randomized natural field experiment (Harrison & List, 2004) on one of Germany's largest massive open online course (MOOC) platforms, openHPI. This nonprofit online education platform provides a public good partly using personal data as inputs: free online education that can be tailored to individual-specific needs. Our intervention aims at increasing the quantity (both at the extensive and intensive margins) and the diversity of personal data available to the MOOC platform – our exemplary public good provider – such that the platform can supply the best fitting services to all learners.

Our experiment compares one control group to two treatment groups. The control group receives a pop-up message which prompts users to complete their profile and hence draws their attention to their profiles. The treatment conditions go beyond this pure attention effect. In the first treatment (*Public Benefit*), the pop-up message additionally makes the public benefit of sharing personal data on the platform more salient (Chetty *et al.*, 2009; Bordalo *et al.*, 2013), thereby emphasizing that providing information has positive effects beyond private benefits. Resembling Andreoni (2007) and Zhang and Zhu (2011) raises the perceived number of beneficiaries that benefit from contributing personal data, which may trigger more contributions. In the second treatment (*Public Benefit + Privacy*), the pop-up message additionally highlights data protection standards, thereby reducing potentially overestimated privacy costs. We measure the completeness of user profiles before and after the pop-up intervention. This experimental design allows us to investigate whether classical interventions, which mitigate the underprovision of public good contributions, increase the amount of personal data donated and whether it is necessary to additionally account for their privacy-sensitive nature. Furthermore, by comparing profile content, the experimental design also allows to study how the treatment affects the diversity of the contribution base.

Overall, our treatments increase personal data contributions at the intensive and extensive margins, and make the database more diverse. At baseline, 48% of users have an entry in their profile and the average user completes 2.6 (out of 11) profile categories. Making public benefits more salient significantly boosts profile completeness by 5.3% compared to the control group. If combined with an emphasis on privacy protection, this effect increases further, but not significantly further, to 6.4%. These effects are sizable⁶ given that we observe a higher intention to update

⁶While our effect sizes are substantially smaller than those in Athey *et al.* (2017), who report 50% less reluctance to share friends' correct email addresses in exchange for a pizza, we start from a very different

one's profile in the control group, where after seeing the pop-up more users click on the link to their profile than in the treatment groups. While we find no clear evidence for treatment effects on the overall extensive margin, that is, whether users have at least one profile entry, we do observe an increase by 12% on the extensive margin when examining the four most privacy-sensitive categories. Furthermore, the type of users who contribute their personal data changes significantly, especially in the *Public Benefit + Privacy* treatment. For instance, after the intervention, user characteristics shift in terms of job experience, job position, education, and gender, generating a more diverse database. We particularly observe such shifts in the distribution of user characteristics for the more sensitive personal information. These results imply that internalizing public benefits increases personal data contributions to public goods and that accounting for the privacy-sensitive nature of personal data tends to make mitigating underprovision more effective, especially when it comes to more sensitive and diverse personal data contributions.

Our article relates and contributes to the literature on mitigating the underprovision of public goods in two ways. First, we gauge whether using insights on how to increase individual-specific contributions to public goods apply to personal data as contributions as well. Previous research studies feedback giving and knowledge sharing as forms of individual-specific contributions to public goods. Results by Cabral and Li (2015) suggest that the underprovision of feedback cannot successfully be tackled with monetary incentives. In contrast, behaviorally motivated interventions appear more successful in mitigating underprovision. For instance, reputation systems and social comparisons can increase feedback provision (Bolton *et al.*, 2004; Chen *et al.*, 2010).⁷ Yet, as results by Ling *et al.* (2005) show the exact wording is important for actually achieving positive effects.⁸ For knowledge as an individual-specific contribution to a public good, studies on Wikipedia show that a combination of private and public benefits (Chen *et al.*, 2020) as well as a large number of beneficiaries (Zhang & Zhu, 2011) determine contributions to the public information good.⁹ Building on these insights, we implement behaviorally informed interventions in a field experiment, which aim at increasing personal data sharing, a new form of an individual-specific contribution to a public good. In particular, we increase the

baseline [2.6 out of 11 answers, i.e., 23.6% relative to 5% in Athey *et al.* (2017)] and use a weaker non-monetary incentive. In a survey experiment by Marreiros *et al.* (2017), privacy salience interventions decrease disclosure of name and email address by 20–30% when being informed in the study description that the study is about online privacy. Hence, given our much more subtle interventions, we consider our effect sizes as meaningful.

⁷With larger social distance, which may be particularly relevant on online platforms, the underprovision of feedback worsens (Bolton *et al.*, 2020).

⁸In Ling *et al.* (2005), emphasizing the uniqueness of one's individual-specific feedback increased contributions, while highlighting public and personal benefits had the opposite effect.

⁹With respect to laboratory evidence on the relationship between group size and public good provision, early studies, as reviewed by Ledyard (1995), find ambiguous results. In contrast, Andreoni (2007) reports that doubling the number of beneficiaries increases contributions but not by the same amount. Diederich *et al.* (2016) find a positive effect of group size in a linear public good game with a large, heterogeneous subject pool. Goeree *et al.* (2002) also estimate a positive relationship. Wang & Zudenkova (2016) claim that there is a discontinuous relationship between contributions to public goods and group size with the relationship being positive for small groups.

salience of the public benefit when contributing personal information to a public education good and can show that this increases contributions.

Second, we contribute to research in the domain of privacy by investigating the effect that privacy sensitivity of personal data has on data provision. Research in this domain has so far focused on pricing or sharing personal data under varying data protection standards in other than public benefit-enhancing settings.¹⁰ While in our setting personal data are not sold for profit but serve the common good, this literature strand provides guidance for our experimental design. For one thing, it suggests that contextual cues affect the sharing of personal information. For example, there are differences in personal data sharing based on which heading a personal data survey has, whether privacy rating icons are displayed, or at which position the privacy-protecting items are listed (John *et al.*, 2011; Tsai *et al.*, 2011; Chang *et al.*, 2016; Athey *et al.*, 2017). For another, it shows that salience rather than the actual comprehensiveness of privacy protection appears important when individuals decide about sharing personal information (Tucker, 2014; Athey *et al.*, 2017; Marreiros *et al.*, 2017).¹¹ Furthermore, data sharing has privacy costs, which may increase when disclosure is incentivized (Ackfeld & Güth, 2019). Therefore, our experiment not only makes the public benefit more salient but also data protection. Our results highlight that personal data contributions to a public good increase the most relative to baseline when privacy concerns are accounted for, particularly for sensitive information and from heterogeneous contributors.

The remainder of this article is structured as follows. The ‘Experimental setup’ section describes the data and the experimental design. Our empirical strategy is outlined in the ‘Empirical strategy’ section. The ‘Results’ section presents the experimental results, and the ‘Conclusion’ section concludes.

Experimental setup

Online platform environment

We conduct our field experiment on one of the biggest German MOOC platforms with more than 200,000 users, openHPI, which offers free online courses covering topics in computer science as well as information and communication technology for beginners and experts either in English or German. We implement our experiment in four courses offered between September 2019 and February 2020, namely ‘Network virtualization – from simple to cloud,’ ‘Introduction to successful remote teamwork,’ ‘The technology that has changed the world – 50 years of internet,’

¹⁰Regrading pricing privacy, Tsai *et al.* (2011), Beresford *et al.* (2012), Jentzsch *et al.* (2012), and Benndorf and Normann (2018) try to elicit a monetary value of privacy. Feri *et al.* (2016) show that some subjects react to the risk of privacy breaches.

¹¹When confronted with information about online companies’ privacy policies, subjects in Marreiros *et al.* (2017) are less willing to share personal information independent of whether the information regarding companies’ privacy protection standards are positive or negative. In contrast, Athey *et al.* (2017) find that learning about an irrelevant privacy-enhancing encryption technology reduces rather than increases the desire to protect one’s privacy. With respect to privacy in advertising, Tucker (2014) shows that shifting the perception but not the actual control over personal profile information on Facebook raises the willingness to click on personalized ads. While the data (Tucker, 2014) use stem from an awareness campaign of a non-profit organization, the intervention is still used to generate higher revenues for an external party.

and ‘Data engineering and data science.’¹² While slightly different in structure, all courses have the same enrollment procedures and consist of video lectures and individual or group assignments. Moreover, all courses use the same interface and have the same requirements to earn certificates.¹³

Our intervention targets the user profile. To enroll in a course, one must register as a user on the openHPI platform providing a valid email address for communication and a (real) name that will be printed on course certificates. During registration, a user profile is automatically created and can be updated by users at any time. Besides these required fields, users can voluntarily provide the following information in their profiles which are not visible to other users: date of birth, company affiliation, career status, highest educational degree, professional experience, professional position, city, gender, country, main motivation for taking courses on the platform, and regular computer usage.¹⁴ The last two profile categories were introduced just shortly prior to our intervention, and all other categories had been part of the profile at the time of registry of all users.¹⁵ We use the new profile categories to rationalize the appearance of our intervention’s pop-up message in courses.¹⁶

Hypotheses and experimental design

In order to test what fosters personal data contributions to a digital public good, we investigate two factors. First, we test a rather classical remedy to underprovision. We follow results by Andreoni (2007) and Zhang & Zhu (2011) in the domains of monetary and knowledge contributions to public goods, who show that contributions to public goods increase with the number of beneficiaries. Following this insight, making public benefits of personal data contributions salient (Chetty *et al.*, 2009; Bordalo *et al.*, 2013) by raising the perception of the number of beneficiaries on the online learning platform should have a positive effect on contributions. To test whether this insight transfers to the domain of personal data contributions, we formulate the following hypothesis:

Hypothesis 1: Emphasizing the *Public Benefit* of contribution increases personal data contributions relative to a *Control* message.

¹²We discuss how our specific sample may affect results in the conclusion.

¹³In most courses, participants can earn a ‘Confirmation of Participation’ if accessing 50% of the material. When achieving 50% of points in the assignments and the final exam, participants receive a ‘Record of Achievement.’ While the course material can also be accessed after the scheduled course dates, graded exercises and tests are no longer available afterwards.

¹⁴Additionally, users can define a display name as a pseudonym, which is used in the course forum. However, this does not contain any relevant, real-world information about users and is therefore disregarded in our analysis.

¹⁵The new computer use category replaced one that was also related to computer usage but contained less distinguishable inputs in terms of content.

¹⁶The new categories were published 8 days before the start of the first treated course. Hence, it is unlikely that participants have encountered the new categories before being directed to their user profiles via our interventions. Only 5.8% of users in our sample updated their profiles independently before seeing the intervention pop-up. To control for these updates in the analysis, pre-intervention profile entries are included as a covariate. The different time spans between the publishing date and a course’s starting date are captured by course dummies in the regressions.

Table 1. Wording for control and treatment groups.

Control	Public Benefit	Public Benefit + Privacy
Dear Learner, We have updated our profile categories. Please take a moment to complete your profile.	Dear Learner, We have updated our profile categories. Please take a moment to complete your profile.	Dear Learner, We have updated our profile categories. Please take a moment to complete your profile.
	By providing your information, you support openHPI in improving its free online education services and the learning experience for the whole openHPI community.	By providing your information, you support openHPI in improving its free online education services and the learning experience for the whole openHPI community.
		Your profile will only be visible to you and the openHPI team but not to other openHPI users. Your data will only be used for research and platform improvement in accordance with our data protection standards.

Notes: The words ‘data protection’ in *Public Benefit + Privacy* contain a link to the privacy protection guidelines of the platform. All treatments include a link to the user profile at the end. Supplementary Figure A.1 provides a screenshot of the pop-up.

For the second hypothesis, we take into account that personal data by its very nature may be privacy-sensitive. These privacy concerns may attenuate the positive effect of highlighting the number of beneficiaries. If this is true, we expect personal data contributions to rise more strongly if greater salience of public benefit is combined with making privacy protection salient. This may reduce potentially overestimated privacy costs of data sharing.

Hypothesis 2: Additionally emphasizing data protection standards further increases personal data contributions relative to just highlighting the *Public Benefit*.

Users who actively engage with the MOOC material after the first week of the course are randomly assigned to a control group and two treatment groups. Randomization is implemented based on the standard round-robin algorithm to perform circular assignments ensuring equal group sizes.¹⁷ Thereby, no stratification is applied. If a user is enrolled in more than one course, we only count her in the chronologically first course. However, it does not matter whether the user was already enrolled in other courses on the platform that are not part of the experiment. Table 1 shows the different treatment texts, Supplementary Figure A.1 displays a screenshot of the most comprehensive pop-up, and Figure 1 graphically represents the timeline of the experiment.

¹⁷Due to technical reasons, we have to exclude users who access the course material exclusively via the mobile app. This reduces the sample size by 509 potential observations to 6155.

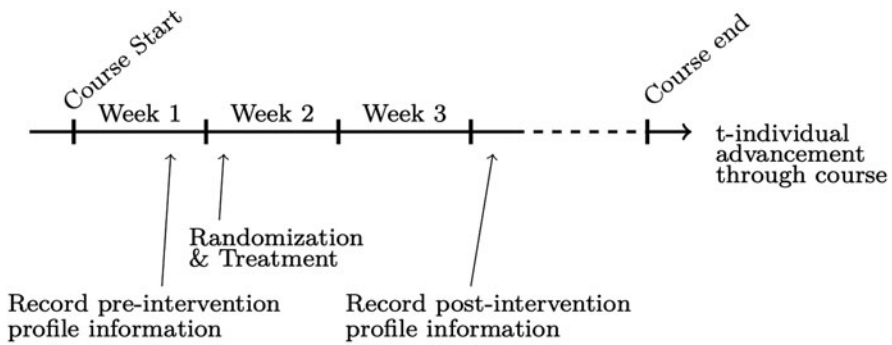


Figure 1. Timeline of the experiment.

Our sample includes all users who are active in the second course week; hence, by design, we exclude users from our sample and intervention that do not make it that far.¹⁸ Two aspects guided this decision. First, the first week is already full of information; hence, participants could miss our intervention or important course-related information.¹⁹ Second, platform improvement based on extended user information is meant to target users with a genuine interest in courses. By focusing on participants who are still active in the second week, we exclude only marginally interested participants.

For the pre-intervention baseline, we record profile completeness 2 days before the intervention (Days 5–6 of the course). More precisely, we measure whether users have any profile entries and if so, how many profile entries. We compare this with the profile completeness 21–22 days after course start, that is, 14–15 days after our intervention. This gives course participants 2 weeks to edit their profiles in response to the intervention. Collecting post-intervention data after 2 weeks allows us to also include those users in our sample, who lagged behind at the beginning of the course but caught up in between.²⁰

The pop-up messages in the two treatments and the *Control* group all contain the following text and a link to the user's profile: *'Dear Learner, We have updated our profile categories. Please take a moment to complete your profile.'* The *Control* group pop-up ensures that we can isolate a pure reminder effect of the pop-up message from effects due to the salience of public benefits and overestimated privacy costs which are the focus of this article. In the *Public Benefit* treatment, the standard text is extended by a note on the public benefit that providing personal information can have for the whole user community. It reads: *'By providing your information, you support openHPI in improving its free online education services and the learning experience for the whole openHPI community.'*

¹⁸Approximately one third of enrolled course participants reaches the second course week's material.

¹⁹For example, in the 'International Teams' course, there is a planning prompt pop-up when participants access the course material the first time, which would stand in conflict with our treatment pop-up.

²⁰11.0% of users in our sample access the course for the first time more than 7 days after course start.

In the *Public Benefit + Privacy* treatment, a remark is added to this statement emphasizing privacy protection standards, particularly who has access to the shared information: ‘Your profile will only be visible to you and the openHPI team but not to other openHPI users. Your data will only be used for research and platform improvement in accordance with our data protection standards.’ The reference to data protection includes a link to the data protection webpage.²¹

Descriptive statistics at baseline

This section reports descriptive statistics of our baseline pre-intervention sample and additionally allows to check whether randomization into treatments was successful. First, we document the pre-intervention outcomes for all treatment groups in Panel A of [Table 2](#). 48.0% of users have at least one entry in their profile before the intervention, and the average profile includes 2.6 completed entries out of 11. There are four categories that treatment-blind raters categorized as disproportionately sensitive: one’s company name, the highest educational degree, professional experience, and the current job position.²² For these categories, we observe much lower baseline values; 35.3% for the extensive and 0.9 for the intensive margin. Before the intervention, most profile categories have a missing rate of at least 60.4% (Supplementary Table A.4). For the two newly introduced categories ‘main motivation’ and ‘regular computer use,’ the missing share is much higher, that is, 94.9% and 94.7%, respectively. χ^2 -tests do not detect any statistically significant differences across the treatment groups in terms of the share of missing values pre-intervention (all $p \geq 0.128$).²³

Second, we report the pre-intervention sample composition in Panel B of [Table 2](#). For 18.9% of users in our sample, the course is the first course they take on the platform. 87.0% access the course from a browser located in Germany. This high share is not surprising given that three out of four courses in our sample are taught in German. 57.3% of users participate in the course ‘Data Engineering & Data Science,’ 18.8% in ‘50 Years of Internet,’ and 19.2% in ‘Network Virtualization.’ Only 5.2% of users participate in the English-speaking course ‘International Teams.’

Third, Panel C of [Table 2](#) describes users’ pre-intervention course behavior and related course information, and confirms that users across treatments are similar in these domains. On average, users enroll 53.7 days prior to course start and begin working on the material 2.7 days after the course start. Since our sample only includes

²¹The link opens a new browser tab with the data protection guidelines. Opening another tab implies that we distract users in the *Public Benefit + Privacy* treatment from editing their entries on the profile page. This diminishes the chance to find a treatment effect in *Public Benefit + Privacy* and thus strengthens any findings.

²²Rating was on a scale from 1 = ‘not at all privacy-sensitive’ to 7 = ‘totally privacy-sensitive.’ Raters were informed that they are rating user profile categories from an online education platform and that these data are only shared with the platform but not with other users, exactly as it is the case on the platform. We sum up the ratings for each category and calculate the average privacy sensitivity. We call a category privacy sensitive if its mean rating is higher than the mean rating over all categories.

²³The entry with the largest difference between treatments is company affiliation ($p = 0.128$). However, this is not surprising given that users in our sample report 563 different affiliations. All other differences are insignificant with a p-value of at least 0.240.

Table 2. Pre-intervention course activity and characteristics overall and by treatment.

	Pooled	Control	Public Benefit	Public Benefit + Privacy	p-value
Panel A: Pre-intervention profile status					
Share of users with at least one entry (extensive margin)	0.480 (0.500)	0.480 (0.500)	0.478 (0.500)	0.481 (0.500)	0.974
- for sensitive categories	0.353 (0.478)	0.358 (0.480)	0.353 (0.478)	0.349 (0.477)	0.642
Number of profile entries per users (intensive margin)	2.560 (3.284)	2.604 (3.304)	2.525 (3.282)	2.542 (3.265)	0.924
- for sensitive categories	0.957 (1.404)	0.985 (1.428)	0.947 (1.397)	0.939 (1.388)	0.811
Panel B: Pre-intervention sample composition					
Course is 1st course	0.189 (0.392)	0.184 (0.388)	0.189 (0.391)	0.1958 (0.396)	0.686
Course accessed from Germany	0.870 (0.336)	0.875 (0.331)	0.870 (0.336)	0.865 (0.341)	0.672
Course 'Network Virtualization'	0.188 (0.391)	0.189 (0.392)	0.189 (0.391)	0.186 (0.389)	0.948
Course 'International Teams'	0.052 (0.221)	0.051 (0.220)	0.050 (0.218)	0.053 (0.225)	0.886
Course '50 Years of Internet'	0.188 (0.391)	0.184 (0.387)	0.188 (0.391)	0.192 (0.394)	0.778
Course 'Data Science and Engineering'	0.573 (0.495)	0.576 (0.494)	0.573 (0.495)	0.569 (0.495)	0.894
Panel C: Pre-intervention course behavior					
Days enrolled after start	-53.7 (83.9)	-53.2 (84.5)	-56.3 (86.0)	-51.7 (81.0)	0.111
Days until first action in course	2.7 (3.9)	2.7 (3.9)	2.6 (3.9)	2.6 (4.0)	0.701
% material accessed in the first week	0.923 (0.207)	0.922 (0.207)	0.926 (0.202)	0.921 (0.211)	0.460
% self-test solved in the first week	0.821 (0.308)	0.823 (0.304)	0.8263 (0.302)	0.815 (0.319)	0.666
N	6155	2052	2060	2043	

Notes: Mean values reported with standard deviations in parentheses. Participants enrolled in more than one intervention course only enter the sample once with their chronologically first course. p-values stem from χ^2 - tests of independence of frequencies between treatments.

users who are still active in the second course week, we observe a high level of first-week activity: users access 92.3% of the material and complete 82.1% of all self-tests in the first course week.

In sum, for all pre-intervention characteristics, we find no statistically or economically significant differences between treatments. All p-values from χ^2 -tests for equal distribution over all treatments exceed the 10% significance level. Thus, randomization into treatment was successful. Furthermore, with more than two thousand observations in each treatment group, we have enough power to identify a 10% effect size at the extensive margin and a 5% effect size at the intensive margin.

Empirical strategy

We estimate the effects of our treatment dummies on post-intervention information disclosure, controlling for the initial disclosure level:

$$y_{it} = \beta_0 + T' \beta_1 + \beta_2 y_{it-1} + X' \beta_3 + \epsilon_{it} \quad (1)$$

for individual i with T' being a vector of treatment dummies, $T_1 = \text{Benefit}$ and $T_2 = \text{Public Benefit} + \text{Privacy}$. We include a matrix (X) of control variables to increase the precision of our estimates. y_{it-1} are the pretreatment outcomes.

As main dependent variables y_{it} , we focus on (1) the extensive margin, that is, whether at least one profile category is filled after the treatment intervention, (2) the intensive margin, that is, how many profile categories are filled, and (3) whether users click on the link to their profile. Clicking on the profile link in the pop-up corresponds to an intention to provide personal data in our experiment. Because there is no baseline for clicking on the link, Equation (1) simplifies to

$$y_i = \beta_0 + T' \beta_1 + X' \beta_2 + \epsilon_i. \quad (2)$$

As secondary outcomes, we look at sensitive categories separately. Furthermore, we study the type of profile changes. Do users only add to their profile? Or, do they also delete and update entries? Updating categories may be relevant if, for example, IT proficiency or work experience has increased since the last revision of the profile.

As controls X , we include several context-related variables into our regressions.²⁴ First, we add course fixed-effects. These dummies do not only capture differences between courses but also different durations between the respective course start date and the publishing of the new profiles categories. Second, we use the enrollment date and the first show-up after course start to control for self-organization skills and the level of course commitment. The latter enters our estimation equation as a dummy variable indicating whether the person accessed the course at least as early as the median user. Third, we include a dummy variable for whether it is the first course the user takes on the platform. This accounts for experience with and potential

²⁴We preregistered two more control variables: a dummy for whether a user allows web tracking and a dummy for whether the user clicks on the link to privacy protection guidelines in the pop-up. However, web tracking was not recorded correctly in all courses and only three users clicked on the privacy protection link, so we refrain from including these controls.

trust toward the platform. Fourth, we control for different reactions between different cultures, for example, with respect to privacy concerns (Bellman *et al.*, 2004; IBM, 2018) by including a dummy variable for course access from Germany, an information that the browser provides.

Results

Main results

In this section, we investigate treatment effects on our three main outcomes of interest: (1) the extensive margin, that is, the share of users that have at least one profile entry after the intervention, (2) the intensive margin, that is, the mean number of profile entries users have after the intervention, and (3) the intention of data sharing, that is, whether users click on the link to their profile in the pop-up.

At the extensive margin, we do not find statistically significant differences between the treatment groups and the control group despite point estimates being twice as large for the *Public Benefit + Privacy* than for the *Public Benefit* treatment group (Table 3, columns 1 and 2; $p = 0.615$ for the *Public Benefit* and $p = 0.250$ for *Public Benefit + Privacy* treatment). Nevertheless, the confidence intervals include effect sizes that are of economic significance. For the *Public Benefit* and the *Public Benefit + Privacy* treatment, the 95% confidence intervals can only rule out effect sizes beyond $(-5.2\%; 8.8\%)$ and $(-2.7\%; 11.3\%)$ relative to the *Control* group, respectively.²⁵ This suggests that our treatment effects are not clear zero effects, but rather imprecisely estimated because we lack statistical power for identifying effects at the extensive margin that are smaller than 10%.²⁶ Thus, at the extensive margin, there is no robust statistical evidence for our hypotheses; the point estimates merely suggest that highlighting public benefits and privacy protection leads to more disclosure than just calling attention to the profile as in the *Control* group.

At the intensive margin, we detect a substantial and statistically significant increase in the amount of profile entries. As Figure 2 shows, this increase differs significantly between the control and treatment groups. In line with our hypotheses, the increases in profile entries in *Public Benefit* and *Public Benefit + Privacy* are statistically significantly larger than in the *Control* group ($p = 0.017$ and $p = 0.005$, *t*-test).²⁷ The largest increase takes place in the *Benefit + Privacy* treatment group, in which the mean of completed profile entries rises from 2.5 to 4.0. Controlling for pre-intervention profile completeness in an OLS regression (Table 3, columns 3 and 4 including further controls), we obtain positive point estimates for both the *Public Benefit* and *Public Benefit + Privacy* treatment indicators which are significant at the 5% and 1% level, respectively. Precisely, users in *Public Benefit* contribute on average 0.18 additional profile entries independent of their pre-intervention profile status compared to

²⁵To illustrate, the effect size confidence interval for the public benefit treatment on having at least one entry was calculated by: $[(0.005 - 1.96 \cdot 0.01)/(0.28); (0.005 + 1.96 \cdot 0.01)/(0.28)]$.

²⁶Given reasonable sample sizes the online platform could provide, *ex ante* we estimated to have enough statistical power to identify effect sizes of 10% or more at the extensive margin.

²⁷The same conclusion holds if only inspecting profile entries which were already part of the profile in the past, that is, all entries except motivation to take courses on the platform and computer usage, as Supplementary Figure A.2 shows ($p = 0.030$ and $p = 0.014$, *t*-test).

Table 3. OLS regression results of main outcomes on treatment.

	At least one entry		Number of entries		Link clicked	
	(1)	(2)	(3)	(4)	(5)	(6)
Public Benefit	0.005 (0.010)	0.005 (0.010)	0.184** (0.080)	0.186** (0.080)	-0.052*** (0.014)	-0.052*** (0.014)
Public Benefit + Privacy	0.012 (0.010)	0.012 (0.010)	0.223*** (0.081)	0.224*** (0.081)	-0.044*** (0.014)	-0.044*** (0.014)
At least one entry (pre-intervention)	0.715*** (0.008)	0.715*** (0.009)				
Number of entries (pre-intervention)			0.878*** (0.008)	0.879*** (0.008)		
Constant	0.280*** (0.010)	0.258*** (0.018)	3.469*** (0.056)	3.210*** (0.134)	0.722*** (0.010)	0.565*** (0.025)
Controls	No	Yes	No	Yes	No	Yes
<i>N</i>	6155	6155	6155	6155	6155	6155
<i>R</i> ²	0.55	0.55	0.55	0.55	0.00	0.02

Notes: Robust standard errors are given in parentheses. Estimations for columns 1–4 as specified in Equation (1) and for columns 5 and 6 as in Equation (2). Controls include dummies for the courses ‘International Teams,’ ‘50 Years of Internet,’ and ‘Data Science & Engineering,’ whether the course is the first course on the platform, whether the course is accessed from Germany, and whether it is accessed earlier than the median access, as well as the day of enrollment relative to the course start. ‘Entries pre’ in columns 3 and 4 are demeaned. This way the constant can be interpreted as the mean effect observed in the *Control* group. We highlight these coefficients with a * symbol.

p* < 0.10, *p* < 0.05, ****p* < 0.01.




Figure 2. Intensive margin: increase in number profile entries by treatment. *Notes:* The figure displays the post-intervention increase in the number of profile entries completed. The error bars indicate the 95% confidence intervals.

Control group users. In *Public Benefit + Privacy*, users even provide 0.22 additional entries. In other words, every fifth treated user in *Public Benefit* and *Public Benefit + Privacy* fills out one more empty profile category than the *Control* group participants do after seeing the pop-up message. The effect sizes amount to 5.3% and 6.4%, respectively.²⁸ While the two treatment coefficients are not statistically different from each other ($p = 0.653$), the larger increase in *Public Benefit + Privacy* suggests that emphasizing privacy protection along with the public benefit may render more personal data contributions. In short, we find evidence for Hypothesis 1 that the salience of public benefits can encourage users to disclose more personal data than when user are just reminded of their profile. Albeit lacking statistical significance, the point estimates also provide suggestive directional evidence for Hypothesis 2. This underlines that positive effects of highlighting the public benefits may be attenuated if privacy concerns are not taken into account.²⁹

Our third main outcome is the intention to update the profile. We approximate intention with the share of clicks on the link to the profile in the intervention pop-up. Surprisingly, significantly more users click on the link in the *Control* group than in the *Public Benefit* and *Public Benefit + Privacy* treatments as [Figure 3](#) displays ($p < 0.001$ and $p = 0.002$, rank sum test). Concretely, 72.2% of *Control* group users click on the link, while 67.0% and 67.7% do so in *Public*

²⁸These effect sizes are calculated as follows: $\frac{\hat{\beta}_{treated}}{\text{Control group mean}} = \frac{0.184}{3.469} = 0.053$ and $\frac{0.223}{3.469} = 0.064$.

²⁹In Supplementary Table A.3 and the discussion thereof, we report further results on the intensive margin with respect to heterogeneous reactions to treatments by different user subgroups (first time users, early course activity, home country, and number of prior profile entries). We do not find any statistically significant heterogeneous effects and therefore conclude that any pop-up message attracts information from all subgroups alike.




Figure 3. Intention to share data: participants who click on profile link. *Notes:* The figure displays the share of users who click on the link to their profile which was included in the pop-up. The error bars indicate the 95% confidence intervals.

Benefit and *Public Benefit + Privacy*, respectively. The same picture prevails if investigating treatment effects in an OLS regression framework (without and with control variables in columns 5 and 6 of Table 3), with treatment effects corresponding to decreases of 7.2% and 6.1%, respectively, in *Public Benefit* and *Public Benefit + Privacy* relative to the baseline. While the higher share of users clicking on the link in the *Control* group is surprising at first glance, it is well in line with convex effort cost of text reading time (Augenblick *et al.*, 2015). The *Control* group text is the shortest, so users may be more likely to read it to the end and are thus more likely to reach the button with the profile link. In light of this potentially offsetting effect, the treatment effects on the intensive margin regarding actual profile filling appear even more sizeable. Particularly, while the treated groups have lower click rates, they are more likely to follow through with their intention to update their profiles and hence to contribute personal data than users in the *Control* group.

At the extensive margin, the different prompt messages do not affect users' willingness to share information significantly due to imprecise estimates, but the point estimates suggest effects may go in the hypothesized directions. Finally, users in the *Public Benefit* and *Public Benefit + Privacy* treatments follow through with their intention to provide data more often, even though they have lower profile click rates than the *Control* group.

Overall, our main results show that especially at the intensive margin, the type of pop-up message matters. Relative to the number of post-intervention entries in the *Control* group, the *Public Benefit* treatment and the *Public Benefit + Privacy* treatment increase available user information significantly, in line with our hypotheses.³⁰ While exact benchmarking is difficult due to differences in outcomes, context, and exact experimental design, we argue our effect sizes are meaningful given the intervention,

³⁰Notably, we observe a large increase from pre- to post-intervention in the *Control* group. This suggests that a simple reminder message by itself seems very effective to motivate users to provide personal details. However, this result is only based on correlation and reminders are not the focus of our investigation.

albeit at the lower end of the spectrum found in the literature. For instance, they are about half the effect sizes found when raising donations for a public radio station and explicitly mentioning high donations of others (Shang & Croson, 2009). When it comes to data inputs, Athey *et al.* (2017) report an increase of about 50% in disclosure of friends' email addresses when college students are incentivized with free pizza. Yet, with 5% their baseline is much lower than ours with 23.6%. In a related survey experiment, Marreiros *et al.* (2017) find that privacy salience interventions decrease disclosure of name and email address by 20–30% when participants are informed that the study is about online privacy.

Further outcomes: types of changes

There may be different types of profile changes masked by the main outcomes. For one thing, users may react differently to treatments depending on how privacy-sensitive they perceive the profile categories. Therefore, we study treatment effects both on the intensive and extensive margins with regard to sharing of sensitive and insensitive personal information, respectively. For another, the intervention may induce changes in different directions. Hence, we analyze the effect on profile extensions, profile reductions, and updates of profile entries separately.

First, not all personal data are equally sensitive. Therefore, treatment-blind raters assessed the profile categories by their privacy sensitivity as sensitive and insensitive categories.³¹ On the extensive margin, that is, with respect to having any sensitive entry in the profile, we observe a significant increase of 2.3 percentage points in the treatment groups relative to the control group. This means that users in the *Public Benefit* and *Public Benefit + Privacy* treatment groups are 12% more likely to have at least one sensitive entry than users in the *Control* group (Table 4). For insensitive entries, there is no such difference. This means that the *Control* message performs similarly well in motivating users to provide insensitive profile information as the *Public Benefit* and *Public Benefit + Privacy* messages. In contrast, the messages in the *Public Benefit* and *Public Benefit + Privacy* treatments induce more users to contribute sensitive information than in the *Control* group. Hence, while overall there was no effect at the extensive margin, increasing the salience of the public benefit of contributing data does have a positive effect as hypothesized for sensitive information.

We also find significant increases at the intensive margin, both for insensitive and sensitive profile categories in the *Public Benefit* and *Public Benefit + Privacy* treatments compared to *Control* (columns 3 and 4 of Table 4). While point estimates for *Public Benefit* and *Public Benefit + Privacy* look similar ($p = 0.682$ and $p = 0.937$, respectively), the effect sizes for sensitive and insensitive entries differ in magnitude because of different baseline levels. In particular, there are 5.1% and 6.1%

³¹We count an entry as sensitive if the mean rating of a category is higher than the individual mean ratings. The sensitive categories are company affiliation, highest educational degree, professional experience, and job position. The insensitive categories are city and country, age, gender, current career status, motivation for joining the platform, and computer usage. In fact, the categories rated as the most sensitive are also those with the highest number of missing values in the pre-intervention sample. Here, we do not include profile categories with nearly endless many potential outcomes, for example, affiliation, country, and city.

Table 4. OLS regression results disaggregated by the sensitivity of entry.

	Sensitive entries		Sensitive entries	
	Yes	No	Yes	No
	Extensive margin		Intensive margin	
	(1)	(2)	(3)	(4)
Public Benefit	0.023** (0.010)	−0.007 (0.006)	0.060** (0.027)	0.071* (0.038)
Public Benefit + Privacy	0.023** (0.010)	−0.007 (0.006)	0.071** (0.028)	0.074** (0.038)
Pre-intervention level	0.795*** (0.006)	0.918*** (0.005)	0.852*** (0.006)	1.011*** (0.004)
Constant	0.190*** (0.008)	0.087*** (0.007)	1.167*** (0.019)	2.013*** (0.026)
<i>N</i>	6155	6155	6155	6155
<i>R</i> ²	0.58	0.87	0.64	0.77

Notes: The table reports OLS regression results on the extensive (columns 1 and 2) and the intensive margin (columns 3 and 4). Robust standard errors are given in parentheses. The ‘Pre-intervention level’ corresponds to ‘Entries pre’ for the intensive margin, that is, the number of completed entries classified, and to ‘At least one entry’ for the extensive margin. All ‘Entries pre’ are transformed to a mean of zero. This way the constant can be interpreted as the effect observed in the *Control* group. We highlight these coefficients with a ‘ ’ symbol. Supplementary Table A.1 confirms all results including control variables.

p* < 0.10, *p* < 0.05, ****p* < 0.01.

Table 5. OLS regression results for type of changes by treatment.

	Type of changes		
	Extensions	Deletions	Updates
	Intensive margin		
	(1)	(2)	(3)
Public Benefit	0.185** (0.080)	0.000 (0.003)	0.008 (0.009)
Public Benefit + Privacy	0.225*** (0.081)	0.002 (0.003)	0.011 (0.009)
Pre-intervention level	-0.119*** (0.008)	0.003*** (0.001)	0.019*** (0.001)
Constant	1.293**** (0.056)	0.004*** (0.002)	0.035**** (0.006)
<i>N</i>	6155	6155	6155
<i>R</i> ²	0.02	0.01	0.05

Notes: The table reports OLS regression results on the intensive margin disaggregating by types of changes. Robust standard errors are given in parentheses. The 'Pre-intervention level' corresponds to 'Entries pre' for the intensive margin, that is, the number of completed entries, and to 'At least one entry' for the extensive margin. The constant can be interpreted as the effect observed in the *Control* group. We highlight these coefficients with a * symbol. **p* < 0.10, ***p* < 0.05, ****p* < 0.01. Supplementary Table A.1 confirms this result including control variables.

more sensitive entries in *Public Benefit* and *Public Benefit + Privacy* than for the mean user in *Control* after the intervention. The treatments effects for insensitive categories amount to 3.5% and 3.7%.³² This suggests that making the public benefit, especially if combined with reference to privacy protection, more salient increases the willingness to share especially sensitive personal information at the intensive margin. Overall, we find strong evidence for Hypothesis 1, that enlarging the perceived circle of beneficiaries increases contribution levels both on the extensive and intensive margins for sensitive profile categories.

Second, our results are driven by profile extensions. As Table 5 shows, the intervention triggered mostly profile extensions (column 1) but nearly no deletions (column 2) or updates (column 3). This is reassuring because our *Public Benefit + Privacy* could have also increased the awareness of a privacy risk and led users to delete their existing entries. Yet, the treatment indicators in the regression on deletions in column 2 are close to zero and insignificant, and the constant – capturing the change in the control group – is small in magnitude. This means that our intervention does not reduce the available data stock. In contrast, the point estimates in column 1 look very similar to those on the intensive margin in the main analysis (Table 3, column 3). Hence, the interventions only triggered changes in line with the directions of Hypotheses 1 and 2 (see the 'Hypotheses and experimental design' section) and did not have any unintended consequences.

³²Note that less categories are rated as sensitive than insensitive, namely four relative to seven. Hence, even given a higher constant in column 2 than in column 1, there is more scope for improvement in insensitive categories.

Shifts in the distribution of personal characteristics

As we highlight throughout, personal data differ from monetary contributions because they are individual-specific. Hence, depending on who contributes, the diversity in the data stock may differ. Therefore, in this section, we evaluate whether our intervention on the online education platform not only creates a larger but also a more diverse data stock thereby elevating potentially selective reporting of characteristics. For this, we check whether the post-intervention user characteristics are more diverse than those prior to the intervention. Therefore, we now study the distribution of user characteristics before and after the intervention.

For all profile categories, the distributions of pre- and post-intervention personal data content differ.³³ We quantify these distributional shifts via marginal effects from multinomial logit regressions. We exploit the panel structure of our data and introduce an individual dummy variable indicator for the post-intervention period. This indicator captures the shift in the distribution of personal characteristics before and after the intervention. The results are reported in Table 6 for each treatment group separately; marginal effects for the pooled sample can be found in Supplementary Table A.5 along with graphs depicting the shifts in reported characteristics from before to after the intervention (Figure A.3).

For work-related characteristics, we see more senior users disclosing information, particularly in the *Public Benefit + Privacy* treatment. For example, for job position, we see a 2.9 percentage point increase in users indicating that they are department heads in the *Public Benefit + Privacy* treatment and shifts away from interns and technicians in this treatment. Moreover, the post-intervention distribution includes more users indicating more than 10 years of work experience rather than between 5 and 10 years compared to the pre-intervention distribution. After the intervention, there tend to be relatively fewer users who report teaching as their profession but more researchers, professionals, or other careers. This effect is again driven by significant distributional shifts in the *Public Benefit + Privacy* treatment.

Focusing on demographics, we observe shifts that point to a more diverse user group than the pre-intervention data suggest. We find a disproportionately strong increase in users reporting a Master-level degree as their highest educational degree. This post-intervention increase is significant in all treatments but particularly pronounced in *Public Benefit + Privacy* with 4.8 percentage points. The shift goes along with a significant decrease in users reporting a 'other' as an educational degree in all treatments, and additionally with significant decreases in users indicating being in high school in the *Public Benefit + Privacy* treatment. Moreover, after the intervention, a higher share of users indicates being female. With an increase of 2.4 percentage points, the shift is most pronounced in the *Public Benefit + Privacy* treatment. Furthermore, we observe more younger users in our sample. Both the shares of users younger than 20 years and that of users in their twenties increase significantly after the intervention, mostly at the expense of users between 40 and 49 years.

For the new profile categories, which elicit motivation for taking courses on the platform and computer proficiency, we see large overall increases in the available

³³We refrain from studying the profile categories country, city, and organizational affiliation because these entries contain too many different realizations as outcomes.

Table 6. Post-intervention shifts in reported characteristics (by treatment).

Outcome	Category	Control		Public Benefit		Public Benefit + Privacy	
		AME	SE	AME	SE	AME	SE
Sensitive categories							
Position	Department head	0.0080	0.009	0.0173	0.011	0.0288**	0.011
	Intern	0.0043	0.004	-0.0027	0.007	-0.0131**	0.005
	Project manager	0.0043	0.011	0.0131	0.013	0.0006	0.015
	Team leader	-0.0032	0.009	-0.0031	0.010	-0.0002	0.010
	Technician	-0.0134	0.012	-0.0247*	0.013	-0.0160	0.014
Professional life	More than 10 years	0.0114	0.010	0.0005	0.010	0.0227**	0.010
	Up to 10 years	-0.0184**	0.008	-0.0085	0.008	-0.0099	0.009
	Up to 5 years	0.0070	0.008	0.0079	0.008	-0.0128	0.008
Highest degree	High school student	-0.0112	0.008	-0.0034	0.008	-0.0173**	0.008
	Bachelor	0.0033	0.007	-0.0021	0.007	0.0025	0.007
	master	0.0351***	0.011	0.0351***	0.011	0.0483***	0.012
	PhD	-0.0007	0.006	0.0087	0.006	0.0073	0.006
	Others	-0.0266***	0.007	-0.0383***	0.008	-0.0408***	0.009
Insensitive categories							
Career status	Academic researcher	0.0027	0.004	0.0004	0.005	0.0066	0.004
	Others	-0.0008	0.007	0.0122	0.008	-0.0020	0.009
	Professional	-0.0006	0.010	-0.0061	0.011	0.0092	0.011
	Student	0.0012	0.006	-0.0018	0.006	-0.0044	0.005
	Teacher	-0.0024	0.004	-0.0048	0.004	-0.0093**	0.004

Gender	Female	0.0197***	0.007	0.0087	0.008	0.0237***	0.008
	Male	-0.0197***	0.007	-0.0087	0.008	-0.0237***	0.008
Age group	<20	0.0034	0.002	0.0010	0.003	0.0060*	0.003
	20–29	0.0090**	0.004	0.0158***	0.006	0.0120**	0.005
	30–39	0.0070	0.006	0.0035	0.006	0.0032	0.006
	40–49	0.0055	0.006	-0.0152**	0.006	-0.0121**	0.006
	50–59	-0.0081	0.006	-0.0088	0.007	-0.0084	0.006
	60–69	-0.0098***	0.004	0.0003	0.005	-0.0033	0.005
	≥70	-0.0070***	0.003	0.0035	0.004	0.0027	0.003

Notes: The table reports average marginal effects from multinomial logit regressions for each profile category on an indicator for the post-intervention period. We combined the more traditional degrees 'Magister' and 'Diplom' with 'Master' because they belong to the same International Standard Classification of Education level. Standard errors are clustered on the user level.

*p < 0.10, **p < 0.05, ***p < 0.01.

information because very few participants had provided this information prior to the intervention. While we find significant shifts in the content for both new categories, we refrain from interpreting these shifts due to the limited number of entries before the intervention. Rather, it is worth noting that most users report a professional motivation (66%) and high or intermediate level of expertise in computer usage (50% and 45%).

Overall, this analysis results in three observations for shifts in the post-intervention distribution of user characteristics: First, shifts are in the same direction for all three groups. Second, they are most pronounced in the *Public Benefit + Privacy* treatment group. Third, most shifts in the distribution affect sensitive profile categories. This means that the *Public Benefit + Privacy* treatment not only increases the amount of data available the most, but is also most effective in generating more diverse personal information donations. Thus, adding an emphasis on privacy-protecting may result in a more adequate estimate of the overall user population. Knowing about this diverse user population, the public good provider may better tailor its services to fit the needs of all users. In the context of online education, this means, for instance, course communication could use gender-specific pronouns, performance dashboards and on-boarding of first-time users could be adapted depending on users' motivation.

Conclusion

In this article, we study how to increase personal data donations to a public good. In a digitized world, such data increasingly serve as inputs to public goods, but – as other types of public good contributions – are likely to be underprovided. However, personal data as contributions face two additional challenges. First, in contrast to money and effort, personal data are individual-specific because each person has different individual characteristics. Therefore, it matters that a diverse range of individuals contribute and not only that a large total amount of contributions is raised. Second, in contrast to other individual-specific contributions like feedback or knowledge, personal data are also privacy-sensitive, which may further bias who contributes. Not respecting these particularities of personal data may lead to biased inputs into algorithms. Hence, it is welfare-enhancing to have a large and representative database as inputs such that the public good provider can accommodate all of its users' needs.

So far it is not clear whether behavioral interventions that have proven effective for other types of contributions to public goods (Frey & Meier, 2004; Ling *et al.*, 2005; Andreoni, 2007; Shang & Croson, 2009; Chen *et al.*, 2010, 2020; Krupka & Croson, 2015) translate to personal data. In a field experiment on one of Germany's largest online education platforms, we show that a classical remedy in the sphere of public goods funding – emphasis on a large circle of beneficiaries (Andreoni, 2007; Zhang & Zhu, 2011) – also significantly increases users' willingness to contribute personal data. Furthermore, we find that the effects of such interventions can be even more pronounced if privacy concerns are additionally accounted for.

Specifically, we find that emphasizing the public benefit of contributing significantly increases personal data contributions. This effect appears more pronounced

if privacy protection is made salient in addition to the pure public benefit, potentially by reducing perceived privacy costs. On the extensive margin, our estimates across all profile categories are imprecisely estimated. However, for more privacy-sensitive categories, we observe significantly more treated users completing at least one entry than control group users. This means making the beneficiary circle and the actual privacy costs salient tends to trigger more disclosure of sensitive information. Furthermore, we find that the distribution of user characteristics after the intervention differs significantly from the pre-intervention distribution. This seems to be especially the case for sensitive characteristics when highlighting both public benefits and privacy protection. Hence, the *Public Benefit + Privacy* treatment not only enlarges but also diversifies the database the most.

While our study unravels new findings on how to increase personal data contributions to a public good, which we argue can be of interest to a more general setting, there are also limitations. First, we study a specific context, that is, online education courses focusing on information technology topics. Hence, our sample consisting of people interested in such topics may differ from the general public and the privacy concerns in our sample may not be representative. On the one hand, online course participants may be more knowledgeable about the usefulness of personal data than the general public, for example, how personal data coupled with data science methods can generate business insights. This means that our sample may be more concerned about privacy than the general public and may have already made very deliberate choices. If we find positive reactions to the salience of privacy protection even in this sample, effects for the general public may be even larger. On the other hand, participants in online courses related to information technology may be less concerned about privacy, for example, because they feel IT-savvy, in which case our sample would have been more nudgeable than the general public. While we are not aware of any study comparing privacy concerns of IT-interested people to the general public, according to a survey by Martin *et al.* (2016), IT professionals care about securing online privacy. Further research may test our privacy salience interventions with a sample that has a less pronounced interest in IT topics.

Second, the opportunity to increase personal data sharing by hinting to privacy protection hinges on the ability of a platform or institution to trustworthily signal privacy protection (Tang *et al.*, 2008; Castro & Bettencourt, 2017; Frik & Mitton, 2019). Without credible data protection in place, the salience of privacy may not increase sharing or may even backfire. Since we study personal data sharing in the context of public goods, we navigate in a context where privacy standards and compliance can be assumed to be very high. This does not only hold for our nonprofit online education platform but also for other public goods, for example, governmental COVID-19 tracing apps or other state-supported services. Our results may be less applicable to settings in which profit-oriented firms try to signal privacy in order to increase their own benefit or to nonprofit organizations which cannot guarantee privacy protection.

Granted these limitations, we see three general takeaways from our results. First, they suggest that the size of beneficiaries positively influences the provision of personal data to a public good similarly to what has been shown in more classical public goods settings (Ledyard, 1995; Goeree *et al.*, 2002; Ling *et al.*, 2005; Andreoni, 2007;

Zhang & Zhu 2011; Diederich *et al.*, 2016; Wang & Zudenkova, 2016; Chen *et al.*, 2020). This means that in the context of a public good that uses personal data as inputs, a simple and inexpensive pop-up message making the public benefit of personal information provision salient can be very effective. This extends prior evidence by Chen *et al.* (2020) to personal data as a new type of contribution.

Second, our results imply that the privacy sensitivity of personal data needs to be taken into account when tackling the underprovision of public goods most effectively. In the treatment, in which we do not only make the public benefit more salient but also the high personal data protection standards, we consistently find the strongest effects compared to the control group. This finding is in contrast to laboratory findings by Marreiros *et al.* (2017) but in line with evidence from illusory privacy protection on Facebook (Tucker, 2014) and evidence by Benndorf *et al.* (2015) and Ackfeld and Güth (2019) that privacy concerns influence personal data sharing in competitive settings. This means that emphasizing privacy protection seems to indeed positively influence personal data sharing for the greater good.

Third, we conclude that reference to public benefits, especially in combination with privacy protection, not only increases available information but also attracts information from a more diverse set of public good contributors. The more diverse and representative this information is, the better the quality of a public good can be. In the online education context we study, a broader database means higher quality inputs to algorithms that help make the learning platform adapt to learners. For example, such algorithms can tailor the educational experience to the individual's learning goals or it can target services such as planning prompts to learners for whom they actually can increase certification rates (Andor *et al.*, 2018). Yet, the usefulness of a broad database goes beyond the online education context. More personal data contributions capturing more diversity can be welfare-enhancing because they allow to train unbiased algorithms and offer public services that fit all users' needs. For instance, the recent COVID-19 pandemic shows that tracing apps can only live up to expectations when many participate truthfully in personal data sharing. Furthermore, a broad database on socio-demographic characteristics – like the profile data we study – can help explain the results of abstract machine learning algorithms and uncover potential inherent biases. Hence, a broader database can contribute to fair and interpretable machine learning in the online education context (Conati *et al.*, 2018; Kizilcec & Lee, *forthcoming*) and beyond.

Supplementary material. To view supplementary material for this article, please visit <https://doi.org/10.1017/bpp.2021.39>.

Acknowledgements. We thank the Hasso Plattner Institute, especially Jan Renz and Christian Willems, for their cooperation. Ackfeld acknowledges funding by the Joachim Herz Foundation via an Add-On Fellowship for Interdisciplinary Economics and by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 741409 – EEC). This project was preregistered at the AEA RCT registry, reference number AEARCTR-0004604, and received ethics approval by the WiSo Ethics Review Board at the University of Cologne, reference number 19023VA. We thank Lisa Bruttel, Marco Caliendo, Matthias Heinz, Lukas Kiessling, Christoph Meinel, Axel Ockenfels, Christoph Schottmüller, Stefan Tübbicke, and participants of Potsdam Research Seminar for Economics for helpful comments. We are also thankful for research assistance by Peter Achmus. All views are the authors' own.

References

- Ackfeld, V. and W. Güth (2019), 'Personal information disclosure under competition for benefits: Is sharing caring?', Working Paper.
- Acquisti, A., C. Taylor and L. Wagman (2016), 'The economics of privacy', *Journal of Economic Literature*, **54**(2): 442–92.
- Andor, M. A., K. M. Fels, J. Renz and S. Rzepka (2018), 'Do planning prompts increase educational success? Evidence from randomized controlled trials in MOOCs', Working Paper.
- Andreoni, J. (2007), 'Giving gifts to groups: How altruism depends on the number of recipients', *Journal of Public Economics*, **91**(9): 1731–49.
- Athey, S., C. Catalini and C. Tucker (2017), 'The digital privacy paradox: Small money, small costs, small talk', NBER Working Paper No. 23488.
- Augenblick, N., M. Niederle and C. Sprenger (2015), 'Working over time: Dynamic inconsistency in real effort tasks', *The Quarterly Journal of Economics*, **130**(3): 1067–115.
- Bellman, S., E. J. Johnson, S. J. Kobrin and G. L. Lohse (2004), 'International differences in information privacy concerns: A global survey of consumers', *The Information Society*, **20**(5): 313–24.
- Benndorf, V. and H.-T. Normann (2018), 'The willingness to sell personal data', *The Scandinavian Journal of Economics*, **120**(4): 1260–78.
- Benndorf, V., D. Kübler and H.-T. Normann (2015), 'Privacy concerns, voluntary disclosure of information, and unraveling: An experiment', *European Economic Review*, **75**: 43–59.
- Beresford, A. R., D. Kübler and S. Preibusch (2012), 'Unwillingness to pay for privacy: A field experiment', *Economics Letters*, **117**(1): 25–7.
- Bolton, G. E., E. Katok and A. Ockenfels (2004), 'How effective are electronic reputation mechanisms? An experimental investigation', *Management Science*, **50**(11): 1587–602.
- Bolton, G. E., J. Mans and A. Ockenfels (2020), 'Norm enforcement in markets: Group identity and the volunteering of feedback', *Economic Journal*, **130**(629): 1248–61.
- Bordalo, P., N. Gennaioli and A. Shleifer (2013), 'Salience and consumer choice', *Journal of Political Economy*, **121**(5): 803–43.
- Cabral, L. and L. (Ivy) Li (2015), 'A dollar for your thoughts: Feedback-conditional rebates on eBay', *Management Science* **61** (9): 2052–63.
- Castro, P. and L. Bettencourt (2017), 'Exploring the predictors and the role of trust and concern in the context of data disclosure to governmental institutions', *Behaviour & Information Technology*, **36**: 321–31.
- Chang, D., E. Krupka, E. Adar and A. Acquisti (2016), 'Engineering information disclosure: Norm shaping designs', in *Proceedings of the 33rd International Conference on Human Factors in Computing Systems*, 587–97.
- Chaudhuri, A. (2011), 'Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature', *Experimental Economics*, **14**(1): 47–83.
- Chen, Y., F. Maxwell Harper, J. Konstan and S. X. Li (2010), 'Social comparisons and contributions to online communities: A field experiment on MovieLens', *American Economic Review*, **100**(4): 1358–98.
- Chen, Y., R. Farzan, R. Kraut, I. YeckehZaare and A. F. Zhang (2020), 'Motivating experts to contribute to digital public goods: A personalized field experiment at Wikipedia', Working Paper.
- Chetty, R., A. Looney and K. Kroft (2009), 'Salience and taxation: Theory and evidence', *American Economic Review*, **99**(4): 1145–77.
- Conati, C., K. Porayska-Pomsta and M. Mavrikis (2018), 'AI in education needs interpretable machine learning: Lessons from open learner modelling', *ArXiv abs/1807.00154*.
- Cramton, P., R. R. Geddes and A. Ockenfels (2019), 'Using technology to eliminate traffic congestion', *Journal of Institutional and Theoretical Economics*, **175**(1): 144–48.
- Diederich, J., T. Goeschl and I. Waichman (2016), 'Group size and the (in)efficiency of pure public good provision', *European Economic Review*, **85**: 272–87.
- Fauvel, S., H. Yu, C. Miao, L. Cui, H. Song, L. Zhang, X. Li and C. Leung (2018), 'Artificial intelligence powered MOOCs: A brief survey', in *2018 IEEE International Conference on Agents (ICA)*, 56–61.
- Feri, F., C. Giannetti and N. Jentsch (2016), 'Disclosure of personal information under risk of privacy shocks', *Journal of Economic Behavior & Organization*, **123**: 138–48.

- Fisman, D. N. and K. B. Laupland (2009), 'The sounds of silence: Public goods, externalities, and the value of infectious disease control programs', *Canadian Journal of Infectious Diseases and Medical Microbiology*, **20**: 39–41.
- Frey, B. S. and S. Meier (2004), 'Social comparisons and pro-social behavior: Testing "conditional cooperation" in a field experiment', *American Economic Review*, **94**(5): 1717–22.
- Frik, A. and L. Mittone (2019), 'Factors influencing the perception of website privacy trustworthiness and users' purchasing intentions: The behavioral economics perspective', *Journal of Theoretical and Applied Electronic Commerce Research*, **14**(3): 89–125.
- Ginsberg, J., M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski and L. Brilliant (2009), 'Detecting influenza epidemics using search engine query data', *Nature*, **457**: 1012–14.
- Goeree, J. K., C. A. Holt and S. K. Laury (2002), 'Private costs and public benefits: unraveling the effects of altruism and noisy behavior', *Journal of Public Economics*, **83**(2): 255–76.
- Goldfarb, A. and C. Tucker (2012), 'Shifts in privacy concerns', *American Economic Review*, **102**(3): 349–53.
- Green, B. (2019), *The smart enough city: Putting technology in its place to reclaim our urban future*, Cambridge, MA: MIT Press.
- Harrison, G. W. and J. A. List (2004), 'Field experiments', *Journal of Economic Literature*, **42**(4): 1009–55.
- IBM (2018), 'The Harris Poll', IBM Cybersecurity, Privacy Research.
- Jentsch, N., S. Preibusch and A. Harasser (2012), 'Study on monetising privacy: An economic model for pricing personal information', European Network and Information Security Agency.
- John, L. K., A. Acquisti and G. Loewenstein (2011), 'Strangers on a plane: Context-dependent willingness to divulge sensitive information', *Journal of Consumer Research*, **37**(5): 858–73.
- Kizilcec, R. F. and H. Lee (forthcoming), 'Algorithmic fairness in education', in W. Holmes and K. Porayska-Pomsta (eds), *Ethics in artificial intelligence in education*, Milton Park: Taylor & Francis.
- Kizilcec, R. F., J. Reich, M. Yeomans, C. Dann, E. Brunskill, G. Lopez, S. Turkay, J. J. Williams and D. Tingley (2020), 'Scaling up behavioral science interventions in online education', *Proceedings of the National Academy of Sciences*, **117**(26): 14900–5.
- Krupka, E. and R. Croson (2015), 'The differential impact of social norms cues on charitable contributions', *Journal of Economic Behavior & Organization*, **128**: 149–58.
- Ledyard, J. O. (1995), 'Public goods: A survey of experimental research', in J. H. Kagel and A. E. Roth (eds), *The Handbook of Experimental Economics*, Princeton, NJ: Princeton University Press, 111–94.
- Ling, K., G. Beenen, P. Ludford, X. Wang, K. Chang, X. Li, D. Cosley, Dan Frankowski, Loren Terveen, Al Mamunur Rashid, Paul Resnick, and Robert Kraut (2005), 'Using social psychology to motivate contributions to online communities', *Journal of Computer-Mediated Communication*, **10**(4): 00–00. doi: 10.1111/j.1083-6101.2005.tb00273.x.
- Lv, Y., Y. Duan, W. Kang, Z. Li and F.-Y. Wang (2014), 'Traffic flow prediction with big data: A deep learning approach', *IEEE Transactions on Intelligent Transportation Systems*, **16**(2): 865–73.
- Marreiros, H., M. Tonin, M. Vlassopoulos and M. C. Schraefel (2017), '"Now that you mention it": A survey experiment on information, inattention and online privacy', *Journal of Economic Behavior & Organization*, **140**: 1–17.
- Martin, N., J. Rice and R. Martin (2016), 'Expectations of privacy and trust: Examining the views of IT professionals', *Behaviour & Information Technology*, **35**(6): 500–10.
- Nature Editorial (2020), 'Show evidence that apps for COVID-19 contact-tracing are secure and effective', *Nature*, **580**(7805): 563–63.
- Obermeyer, Z. and E. J. Emanuel (2016), 'Predicting the future – Big data, machine learning, and clinical medicine', *New England Journal of Medicine*, **375**(13): 1216–19.
- Shang, J. and R. Croson (2009), 'A field experiment in charitable contribution: The impact of social information on the voluntary provision of public goods', *Economic Journal*, **119**(540): 1422–39.
- Tang, Z., Y. (Jeffrey) Hu and M. D. Smith (2008), 'Gaining trust through online privacy protection: Self-regulation, mandatory standards, or caveat emptor', *Journal of Management Information Systems* **24** (4): 153–73.
- Tsai, J. Y., S. Egelman, L. Cranor and A. Acquisti (2011), 'The effect of online privacy information on purchasing behavior: An experimental study', *Information Systems Research*, **22**(2): 254–68.
- Tucker, C. E. (2014), 'Social networks, personalized advertising, and privacy controls', *Journal of Marketing Research*, **51**(5): 546–62.

- Tucker, C. E. (2015), 'Privacy and the internet', in S. P. Anderson, J. Waldfogel and D. Strömberg (eds), *Handbook of media economics*, Amsterdam, Netherlands: North-Holland, Vol. **1B**, 541–62.
- Wang, C. and G. Zudenkova (2016), 'Non-monotonic group-size effect in repeated provision of public goods', *European Economic Review*, **89**: 116–28.
- Yu, H., C. Miao, C. Leung and T. J. White (2017), 'Towards AI-powered personalization in MOOC learning', *npj Science of Learning*, **2**(1): 1–5.
- Zhang, X. (Michael) and F. Zhu (2011), 'Group size and incentives to contribute: A natural experiment at Chinese Wikipedia', *American Economic Review*, **101**(4): 1601–15.

Cite this article: Ackfeld V, Rohloff T, Rzepka S (2021). Increasing personal data contributions for the greater public good: a field experiment on an online education platform. *Behavioural Public Policy* 1–27. <https://doi.org/10.1017/bpp.2021.39>