

NBER WORKING PAPER SERIES

THE DIGITAL PRIVACY PARADOX:
SMALL MONEY, SMALL COSTS, SMALL TALK

Susan Athey
Christian Catalini
Catherine Tucker

Working Paper 23488
<http://www.nber.org/papers/w23488>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2017

We want to thank MIT students Dan Elitzer and Jeremy Rubin for orchestrating the MIT Bitcoin experiment and allowing us to study the results. We also want to thank the MIT Office of the Provost for all their support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w23488.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Susan Athey, Christian Catalini, and Catherine Tucker. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Digital Privacy Paradox: Small Money, Small Costs, Small Talk
Susan Athey, Christian Catalini, and Catherine Tucker
NBER Working Paper No. 23488
June 2017
JEL No. C93,D62,D8,K10,O3,O31,O38

ABSTRACT

'Notice and Choice' has been a mainstay of policies designed to safeguard consumer privacy. This paper investigates distortions in consumer behavior when faced with notice and choice which may limit the ability of consumers to safeguard their privacy using field experiment data from the MIT digital currency experiment. There are three findings. First, the effect small incentives have on disclosure may explain the privacy paradox: Whereas people say they care about privacy, they are willing to relinquish private data quite easily when incentivized to do so. Second, small navigation costs have a tangible effect on how privacy-protective consumers' choices are, often in sharp contrast with individual stated preferences about privacy. Third, the introduction of irrelevant, but reassuring information about privacy protection makes consumers less likely to avoid surveillance, regardless of their stated preferences towards privacy.

Susan Athey
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
athey@stanford.edu

Catherine Tucker
MIT Sloan School of Management
100 Main Street, E62-533
Cambridge, MA 02142
and NBER
cetucker@mit.edu

Christian Catalini
MIT
Sloan School of Management
100 Main Street, E62-480
Cambridge, MA 02142
catalini@mit.edu

1 Introduction

Since the initial formalization of privacy policy towards consumer data in the Privacy Act of 1974, there has been an emphasis on ‘Notice and Choice’ to safeguard privacy. ‘Notice’ gives consumers information about data collection and use, and then consumers make a ‘Choice’ about whether or not to allow their data to be collected or used in that way. These mechanisms may not be sufficient. In this paper, we present evidence about a variety of distortions in the notice and choice process, relating to consumer decisions to share data and choose more or less privacy-protective technologies.

To do this we use data from the Massachusetts Institute of Technology digital currency experiment where every undergraduate student at MIT was offered in the fall of 2014 \$100 in Bitcoin (Catalini and Tucker, 2017). The main focus of the experiment was establishing a cryptocurrency community at MIT. However, as part of the experiment, students had to make at least three digital privacy choices: Whether they wanted to disclose the contact details of their closest friends; whether they wanted to maximize the privacy of their transactions from the public, a commercial intermediary or the government; and whether they subsequently wanted to take additional actions to protect their transaction privacy when using Bitcoin. We use randomized elements of the experiment, often not directly focused on the question of privacy itself, to understand how responsive this demographic is to small changes in incentives, costs and information.

The first randomization offered 50% of the students a small incentive in the form of pizza in exchange for the emails of their friends. The original goal of this randomization was to reconstruct accurate social network information on participants to study Bitcoin adoption on campus. However, it also allows us to examine the effect of small incentives on the willingness to disclose information. The second randomization changed the ordering of wallet technologies and the amount of information given about these wallets. We exploit this

variation to test whether students' choices of wallets were consistent with their stated privacy goals, and whether increasing transparency about the wallets' privacy features improved privacy outcomes. The last randomization was whether or not students were exposed to information about the possibility of using encryption to protect their *initial* disclosure of information to us. Our prior was that this additional information would have increased participants' attention to privacy issues. We investigated whether the additional text made students more likely to protect their privacy from the public by obfuscating their transactions on the Bitcoin public ledger, from the intermediary by not revealing additional identifying information to wallet provider, or from the government by not linking their Bitcoin wallet to a traditional bank account which is subject to government oversight.

We have three main findings. First, the effect small incentives have on disclosure may explain the privacy paradox: Whereas people say they care about privacy, they are willing to relinquish private data quite easily when incentivized to do so. Second, small frictions in navigation costs surrounding privacy choices can have large effects in terms of technology adoption, even in the presence of transparent information about the privacy consequences of those choices. Third, our information treatment on encryption - possibly by giving participants an illusion of protection - surprisingly did not increase privacy-enhancing behavior as we expected, but actually reduced it. After being randomly exposed to irrelevant, but reassuring information about a tangential technology, students were less likely to avoid surveillance in their use of the technology. In all these cases, privacy-decreasing decisions take place regardless of stated preferences for privacy.

This paper contributes to three main literatures.

The first stream is a growing literature that uses field experiments to explore the consequences of psychological, cognitive and social factors on economic choices (Harrison and List, 2004; Gneezy and List, 2006; DellaVigna and Malmendier, 2006; DellaVigna, 2009; DellaVigna et al., 2012; Landry et al., 2010; Allcott and Rogers, 2014). In particular, our results

on the ordering of digital wallets and on the effect of additional transparency on privacy outcomes relate to previous research that has highlighted the role of defaults, salience, and limited attention on individual choices (Madrian and Shea, 2001; Ho and Imai, 2008; Chetty et al., 2009; Bertrand et al., 2010).

The second literature is a policy-oriented literature on notice and consent (Posner, 1981; Acquisti et al., 2016). Work in computer science has tended to suggest that the failure of notice and consent lies with firms who purposely obfuscate their privacy notices (McDonald and Cranor, 2009), and that if presented with transparent information, consumers would make the right privacy choices (Tsai et al., 2011). Echoing this literature, legal commentary on notice and choice has tended to emphasize failures on the part of firms to be sufficiently transparent and compliant with existing policy (Marotta-Wurgler, 2016; Reidenberg et al., 2016). By contrast, our paper suggests that compliance with notice and consent may still not achieve the policy goal of protecting consumer privacy: Due to the effect small frictions can have on shifting consumer behavior away from privacy preferences, consent may not actually reflect true consumer intent.

Another more general contribution of the paper to this literature is to emphasize that practically, the choice about whom to keep data private from is more complex in a world where an individual has to rely on a custodian to store their digital data, whether it be a firm, the government, or an open-source community. This differs from earlier privacy work on notice and consent, which focused on the question of protecting consumer privacy from commercial firms, and where the choice was whether to generate data in the first place.

The third literature is a literature on the privacy paradox. Though the term ‘privacy paradox’ has often been used loosely by policymakers and commentators to cover the general mismatch between stated privacy preferences and behavior, the academic literature has been squarely focused on the disconnect between stated privacy beliefs and behavior on social media websites (Gross and Acquisti, 2005; Barnes, 2006). The most similar paper to ours is

(Adjerid et al., 2013), which shows, in the context of a lab experiment regarding a university social media site, that misdirection encouraged students to volunteer more sensitive academic data. We extend this literature beyond social media, and not only document it using field experimental data, but also show that consumers deviate from their own stated preferences regarding privacy in the presence of small incentives, frictions and irrelevant information.

Though there are policy implications of our paper, it is important to emphasize that our empirical results can be used to support two highly contrasting stances towards privacy protection.

The first policy stance is that our results could be taken as suggesting that consumers' revealed behavior regarding privacy – as revealed by their stated privacy preferences in our surveys– is slanted away from their actual normative preferences (Beshears et al., 2008). This might suggest that consumers need to be protected from themselves, above and beyond the protection given by a notice and choice regime, to ensure that small incentives, search costs or misdirection are not able to slant their choices away from their actual normative preferences.

The second policy stance our results document is that there is a disconnect between stated privacy preferences and revealed preference, but that revealed preference is actually closest to the normative preference. When expressing a preference for privacy is essentially costless as it is in surveys, consumers are eager to express such a preference, but when faced with small costs this taste for privacy quickly dissipates. This would suggest that basing privacy protection on stated preference data regarding privacy expressed in surveys is misguided, especially since such policies have been documented to impose costs on firms (Miller and Tucker, 2011; Kim and Wagman, 2015).

2 Empirical Setting and Data

In the Fall of 2014, the MIT Bitcoin Club raised capital from a group of alumni to give each of 4,494 MIT undergraduates \$100 in Bitcoin. We emphasize that the aim of this experiment was not to study privacy. Instead, the objective of the students was to jumpstart the ecosystem for the digital currency on campus, and expose their peers to the opportunities enabled by cryptocurrencies.¹ However, as part of the signup process, participants were asked for their preferences for privacy, and then subsequently had to make choices regarding what data they provided and how privacy-protecting their subsequent technology choices were. This presents a unique opportunity to explore disconnects between stated privacy behavior and actual privacy choices.

Our data covers the 3,108 undergraduates that signed up for a digital wallet. Participation ranged from 79% among first year students to 62% among fourth year students. International and biology students were slightly less likely to participate (61% and 59% respectively), and enrollment was highest (80%) among electrical engineering and computer science students. About a third of the students in the data (32%) had strong self-assessed programming skills ('Top Coders'), and 55% were male. We complement this survey data that we collected when students signed up for their bitcoin with demographic information about the students provided by the Institutional Research section of the MIT Office of the Provost. Descriptive statistics for our sample are shown in Table 1.

To compare students' stated preferences for the privacy features of digital wallets to their revealed preferences, before students selected a wallet we asked them to rate multiple privacy dimensions as part of the sign up process. In particular, students had to rate wallets in terms of their traceability by peers, the wallet service provider, and the government.²

¹In another paper we evaluate the effect of identity and delay on the diffusion and cashing out process. In terms of use of Bitcoin, by the end of our observation period (February 2016), 13.1% was regularly active, 39% of students had converted their bitcoin back to US dollars ('Cash Out'), and the majority of participants (47.9%) was still holding on to their original bitcoin. (Catalini and Tucker, 2017).

²The survey questions asked how important the privacy features of a digital wallet were on a scale from

We use the students' answers to divide the sample into high versus low taste for privacy from each one of the three audiences.³ According to the answers (see Table 1), 38% of students had high taste for privacy from their peers ('High Privacy from Peers'), 55% from intermediaries ('High Privacy from Intermediary'), and 42% from the government ('High Privacy from Government').

We also build measures of the students' degree of trust in different institutions for financial services in the same way.⁴ Based on the responses, 51% of the sample had high trust in the government ('High Trust in Government'), 26% in a startup ('High Trust in Startup'), and 43% in a retailer ('High Trust in Retailer') to provide services such as digital wallets, credit or debit cards, or mobile payments.

There were three randomizations within the experiment that we focus on since they have implications for privacy. The first is whether being given a small incentive to provide friends' emails changed behavior ('Incentive Randomization', randomly shown to 50% of the sample, see Figure A-1b). The key dependent variable for this part of the analysis is a binary indicator equal to one if the student provided us only with invalid emails for their friends (All Invalid).

The second randomization we use was the random ordering of wallets as a source of exogenous variation in wallet choice. This allows us to look at the propensity of students to select a wallet that maximized privacy on different dimensions as a function of the order in which wallets were presented on the page. Figures A-3a and A-3b are two examples of this randomization.

1 (not at all) to 5 (very important). The dimensions used were: "Trackability of your transactions by the government", "Trackability of your transactions by the service provider", "Trackability of your transactions by your peers". The order the features were listed in was randomized.

³Students who do not answer a specific question are grouped in the high privacy part of the sample, as not answering could be a reflection of their privacy attitude. Results are robust to including them in the opposite group or removing them.

⁴The relevant survey questions asked participants "To what extent do you trust the following entities to provide financial services such as digital wallets, credit or debit cards, or mobile payment services?" - and the scale used was from 1 (not at all) to 5 (to a great extent).

The third randomization we used was whether or not we included an additional text about encryption. This ‘Encryption Randomization’, was randomly shown to 50% of the sample, and is depicted in Figure A-5. The text highlighted how ‘Pretty Good Privacy’ (PGP) software can be used to ensure the security of communications between a sender and a receiver through encryption. We measured responses to this randomization by observing whether or not students took one of three actions to reduce the possibility of subsequent surveillance. The first outcome we use is whether students performed any coin mixing to make it more difficult for the public to trace their transactions on the Bitcoin blockchain. Mixing is the act of pooling one’s transactions together with others, so that inputs and outputs are more difficult to link to a single entity on the Bitcoin public ledger. In the absence of mixing, it is relatively easy to track entities across transactions (Athey et al., 2016). The second measure captures whether students using a bank-like digital wallet such as Circle or Coinbase, also revealed additional, identifying information to the intermediary (either a mobile phone number or their home address). The third measure is a dummy equal to one if participants linked their wallet to a traditional bank account, making it possibly easier for the government to link their Bitcoin transactions to their transactions in fiat-currency.

3 Results

3.1 Small Money

When asked by the National Cyber Security Alliance (NCSA) in a survey,⁵ 60% of consumers said that they would never feel comfortable sharing their list of contacts if asked. In the same survey, information about one’s contacts ranked as the second most private piece of data,

⁵<https://staysafeonline.org/about-us/news/results-of-consumer-data-privacy-survey-reveal-critical-need-for-all-digital-citizens-to-participate-in-data-privacy-day>

right below social security numbers (68% would never share their social security number when asked).

In order to study the diffusion of Bitcoin on campus, we needed information about the participants' social ties. This posed a challenge, as it is difficult to collect accurate social network information without relying on Facebook Connect, an option that was discarded in this context to avoid attrition due to privacy concerns. Aware that simply asking about the email addresses of one's friends would give us poor coverage, we decided to randomly include a question during the signup process for 50% of our sample that incorporated a small incentive to encourage disclosure: A pizza that participants could share with their closest friends. This allows us to compare the choices students made in terms of protecting (or not) the privacy of their friends under the non-incentivized ('Ask')⁶ and the incentivized ('Ask + Incentive') regime.⁷

Our key outcome variable in this section is whether students decided to protect the privacy of their friends by giving us invalid addresses or not. Both in the incentivized and in the non-incentivized regime, our dependent variable is equal to one if students provided all invalid emails, and zero otherwise. Since students could only list MIT addresses during the sign up process, we are able to check the validity of these entries by using the MIT People Directory.⁸ We focus on cases where all emails provided are invalid to rule out typing errors, and identify the subset of students that clearly did not want to share this type of information with us.

In the raw data, within the subsample randomly exposed to the incentive, 5% of students

⁶The non-incentivized question, which was presented to the full sample, used the following text: *"List 5 friends you would like to know the public addresses of. We will email you their addresses if they sign up for the directory."*

⁷"The incentivized question, which was randomly presented to 50% of the sample, used the following text: *"You have been selected for an additional, short survey (1 question). If you decide to complete the survey, you will receive one free pizza that you can share with your friends. List 3 friends you would like to share a pizza with. One pizza will be on us! If you happen to talk about Bitcoin, even better!"*.

⁸Available online at: <http://web.mit.edu/people.html>

gave all invalid emails under the ‘Ask’ condition, and 2.5% under the ‘Ask + Incentive’ condition (see Figure 1a). Within the full sample, 6% of students gave all invalid emails under the ‘Ask’ condition.⁹ Figures 1b, 1c and 1d suggests that there is very little variation in how the incentive affects students with high versus low preference for privacy from the public, intermediary, or government. Instead it appears that whatever the stated privacy preference is, students share their friends’ data when given a small incentive.

Table 2 uses Ordinary Least Squares regressions at the student-answer level to test robustness to the inclusion of additional controls and interactions. Only the 1,543 students that were exposed to both the incentivized and the non-incentivized question about the emails of their friends appear in this sample (two rows for each student, i.e. one row for each answer). All columns use robust standard errors clustered at the student level. The incentivized condition has a large, negative effect on the probability that students will protect the privacy of their friends relative to their behavior in the non-incentivized condition. In Column (1), the coefficient estimate of -0.0285 for ‘Ask + Incentive’ represents a 54% decrease in the probability of all invalid emails over the baseline. In Appendix Table A-1 we test the robustness of this result to a number of alternative explanations. For example, one may worry that the effect is driven by students who do not value the contacts of their friends yet because they are only three months into the program, but we do not find heterogeneous effects by cohort. Differences in gender, expectations about the price of Bitcoin, and technology preferences (e.g. digital wallets, browsers etc.) also do not have a meaningful effect on the impact of the incentive.

The main result is also stable when we add in Table 2 interactions between the main effect (‘Ask + Incentive’) and the students’ stated preferences for privacy across different audiences:

⁹When we explore heterogeneous effects by gender, year of study, digital wallet selected, expectations about Bitcoin, coding skills and technology preferences such as operating system or browser used, we find no significant differences in how these subgroups respond to the incentivized regime in the raw data. In all cases, when the request is made together with the pizza incentive, students are significantly less likely to protect the privacy of their friends.

Privacy from peers (Column 2), from an intermediary (Column 3), and from the government (Column 4). In all cases, consistent with the raw data evidence from Figures 1a to 1d, the main effect is qualitatively unchanged and the interactions are insignificant, suggesting that privacy-sensitive individuals do not respond differently to the incentive compared to other individuals. *Ex ante* stated preferences about privacy, at least in this setting, do not seem to separate students in terms of how they will respond to our two conditions.

The absence of heterogeneous effects on these privacy dimensions is somewhat puzzling. One interpretation is that this particular demographic is comfortable with sharing information because it already enjoys limited digital privacy, and therefore incurs limited costs with additional disclosure. To further explore this possibility, Columns (5) to (7) of Table 2 use the same approach as the previous columns, but rely on our survey measures of trust in different institutions for the provision of financial services which, as before, are split into high versus low groupings. The coefficient for the interaction term for high trust in startups and retailers is positive, although in both cases non-significant. A look at the raw data for the subsample exposed to both regimes suggests that the sign is driven by the fact that these students are somewhat less likely to protect their friends' emails in the first place. Whereas students with low trust in startups on average deliver invalid emails in 5.7% of the cases in the non-incentivized regime, their high trust peers do so in 4% of the cases (1.7% difference, $p = 0.1792$). Similarly, students who trust retailers only protect the emails of their friends in 3.9% of the cases, compared to 6.4% for the rest of the students (2.5% difference, $p = 0.0273$).

The results of this section highlight how small incentives such as a pizza can have a large effect on decisions about privacy. While this first part of the analysis focused on the decision to protect the privacy of one's friends, the next two sections will directly focus on choices that affect the focal individual, and quantify how small frictions and information can shift individuals away from their stated privacy goals.

3.2 Small Costs

During signup, students were presented with four Bitcoin wallets, randomly ordered on the page. The randomized order in which digital wallets were presented to the students allows us to explore if introducing small frictions in a sign up flow can change long-term privacy outcomes. For example, if undue haste or inattention induce students to default to the first listed option and ignore the privacy features of each wallet, then the ranking should have a meaningful effect on the wallet students end up using and the data they end up disclosing.

Whereas open-source Bitcoin wallets like Electrum offer a high degree of privacy from the government and do not require an intermediary to be used, they also record all transactions on the Bitcoin public ledger using pseudonyms (Bitcoin addresses). Though it is in theory possible to make it more difficult to tie addresses to entities, in practice such efforts can be undermined: For example, Athey et al. (2016) use different heuristics and public data sources to map pseudonyms to individual entities, and track individual transaction patterns over time such as trading and speculation, international money transfer, and gambling.

Bank-like wallets, instead, connect to traditional bank accounts and credit cards, offer a mobile app, can easily convert Bitcoin to and from government-issued money, and may provide additional privacy to their users from the public because of the way they pool transactions within their network without recording each one of them on the public ledger. At the same time, with bank-like wallets users need to be comfortable sharing all their transaction data and identity information with a commercial intermediary, and possibly the government since these intermediaries need to comply with Anti-Money Laundering (AML) and Know Your Customer (KYC) regulations like other financial institutions.

Students' wallet choices therefore involve a trade-off in terms of who may have easier access to their financial transaction data in the future. The vast majority of participants (71%) selected a bank-like wallet and only 9% selected a wallet that is more difficult for the

government to track because it does not rely on an intermediary.¹⁰

Choices were strongly affected by the random ordering of wallets: When a bank-like wallet was listed first, 78% of students selected it (as opposed to only 65% when it was listed 2nd or lower); when the open-source Electrum wallet was listed first, 12% of students chose it, compared to only 8% when it was not. Small frictions, such as those generated by the ranking of options on a web page, generated large differences in the technology adopted.

Table 3 reports in regression format the effect of wallet ordering on technology choices. Columns (1) to (3) use an indicator equal to one if the focal student selected a wallet that does not record all transactions to the public Bitcoin blockchain. Similarly, in Column (4) to (6) the dependent variable is equal to one if the chosen wallet does not given an intermediary access to transaction data, and in Column (7) to (9) it is equal to one in cases where students selected an open source wallet that is harder to track for the government. In each OLS regression the key explanatory variable, ‘Best Not 1st’, is a binary indicator equal to one if none of the wallets that would maximize privacy along the focal dimension is listed first. Specifically, the indicator ‘Best Not 1st’ is equal to one when additional costs are introduced in the selection of the optimal wallet for the specific dimension of privacy captured by the dependent variable.

The results highlight how the costs introduced by the random order of wallets shape student choices: In Column (1), when wallets that would maximize privacy from the public are not listed first, students are 13% less likely to select them, which corresponds to a 16.7% decrease relative to the baseline. Adding an interaction between the main effect and the participants’ stated preferences for privacy from peers, similar to what we saw in Table 2, has little effect on privacy choices when small frictions in search costs are introduced.

One explanation for the sizable shifts in privacy outcomes we observe is that they are the result of participants selecting wallets under limited information. To investigate this, we

¹⁰As a comparison, only 12.5% of students were using an open-source browser during registration.

compare the baseline condition to a situation where participants had far more information for making their choices, and see whether such information can compensate for the ranking effects.

A screenshot of the randomized treatment we rely on for this part of the analysis is presented in Figure A-4: Whereas 50% of the sample only saw the names, logos and short descriptions of the four randomly sorted wallets (see Figures A-3a and A-3b for two examples), the remaining 50% received additional information about key privacy, security and convenience trade-offs. In particular, under the ‘Increased Transparency’ ballot screen (Figure A-4), students received information about the wallets’ strengths and weaknesses in terms of privacy from the public and an intermediary (Column 2), data security (Column 3), data recovery (Column 4), ability to convert bitcoin to and from US dollars (Column 5), and privacy from the government (Column 6).

Column (3) of Table 3 introduces the ‘Increased Transparency’ randomization and interacts it with the ordering of wallets on a student’s screen. Relative to the baseline case where the best wallet was listed first and no additional information was provided, as before, a lower ranking induces a substantial drop in the likelihood that the participant will maximize privacy along the focal dimension. If we just look at point estimates, these results suggest that while additional transparency can dilute a small amount of the effect of navigation costs, it does not in any way eradicate them.

In Column (4), when we look at maximizing privacy from an intermediary, the reduction in the probability of selecting an optimal wallet is similar (13.2%), but this time maps to a larger relative effect, as the probability of making this type of choice is otherwise only 28.6%. When a wallet that optimizes privacy from an intermediary is not listed first, we observe a 46% decrease relative to the baseline. As before in Column (5), controlling for the students’ stated preferences for privacy from an intermediary and interacting it with the main effect of the best wallet not being featured first has no effect. Similarly, when we add an interaction

between the presence of increased transparency and the ordering of wallets, the importance of navigation costs persists.

Comparable effects are observed in Column (7), where a 3.8% reduction in the probability of selecting the wallet that maximizes privacy from the government corresponds to a 31.6% change over the case where such a wallet is listed first. Once more, controlling for stated preferences for privacy from the government in Column (8) reinforces our main finding. Last, in Column (6) we again show that even with increased transparency, navigation costs still shape decision making.

Taken together, results from Table 3 support the idea that across three different privacy dimensions, the ordering of wallets on the sign up page, by introducing minor, additional costs for the wallets not ranked first, had a large effect on the technology ultimately adopted by the students. Even in an environment where students could maximize privacy in a way that was consistent with their stated preferences, the ordering – potentially combined with inattention or undue haste – seemed to drive many of the participants’ decisions. Across all three dimensions, students that had high taste for privacy on the focal dimension behave no differently than others. Moreover, providing additional information only partially counteracts the effect of small frictions on digital privacy choices.

3.3 Small Talk

In this last section, we study the impact of a small information treatment that explicitly focused on privacy-protecting behavior. Under the ‘Encryption Randomization’ condition (Figure A-5), when shown details about PGP technology, 50% of the sample was provided with extra information on how PGP allows for secure communication between a sender and a receiver, and reduces the ability of a third-party to intercept it. In particular, the randomization focused on how PGP can help individuals *“keep the prying eyes of everyone from governments to Internet service providers [...] from seeing the content of messages.”*

Whereas 55% of participants initially tried this additional step of adding PGP encryption, only 49% of those who tried succeeded, with the others falling back to the easier flow without encryption. This is consistent with many students caring about privacy and security, but then falling back to the most convenient options when additional effort is required, consistent with Section 3.2.

Though PGP encryption technology is widely used in contexts where security and privacy are paramount, in our setting the technology did not provide the students with any additional protection with respect to their future Bitcoin transactions. These would be still exposed to intermediaries, governments or the public depending on the digital wallet selected by the students. By using PGP to encrypt and sign their wallet address before communicating it to us, students could make sure that if a malevolent actor had intercepted the communication and replaced their address with a different one (e.g. to divert the funds), then the PGP signature would not have matched the student’s public PGP key, allowing us to identify the attack. Therefore, PGP was only used to secure the communication of the address to us for the initial distribution of bitcoin. Nevertheless, students may have interpreted the additional information about PGP as relevant for the privacy of all their future bitcoin transactions, but in this setting it was not.

Table 4 estimates the effect of the ‘Encryption Randomization’ on the likelihood of escaping further surveillance by the public, the intermediary or the government. All columns report OLS regressions, and exclude students who abandon Bitcoin, as the privacy outcomes studied here are not relevant for non-adopters. In Columns (1) and (2), the dependent variable is equal to one if the students used a privacy-enhancing mixing service to increase the anonymity of their bitcoin transactions. Mixing services allow users to pool their transactions (multiple inputs and outputs) to make it substantially more difficult for the public to follow the digital trail recorded on the public Bitcoin blockchain. Users may use such a service if they are worried about the public tracking their spending or trading patterns, or

quantifying their overall bitcoin assets. Since not all users may recognize the advantages of using a mixing service to protect their privacy nor may know exactly how to use it, the baseline is very low: Only 2.3% of students used such a service when not exposed to the ‘Encryption Randomization.’ The percentage goes further down because of the information treatment, which corresponds to a 1% reduction in use, and a 45% decay relative to the baseline. Results are noisy, possibly also because we may not be able to capture some of the most sophisticated methods of transaction mixing or the use of less popular services for doing so. Whereas the 1% reduction in use may seem small, it is important to remember that overall activity with Bitcoin is also low, at 13.1%. In Column (2), when we interact the randomization with the students’ stated privacy preferences, results are qualitatively unchanged as in the previous tables (the main effect for the randomization is insignificant because of the smaller sample size, but comparable in size and standard error to the one in Column 1).

Columns (3) to (6) of Table 4 further limit the sample to ‘Bank-Like’ wallets only, as for these wallets we can observe the students’ decisions to disclose (or not) additional information that may make it easier for the intermediary or the government to track them. The dependent variable in Columns (3) and (4) is equal to one if the students did not reveal their mobile phone number or address to the commercial intermediary, and zero otherwise. The effect of the randomization in Column (3) is negative but small, noisy and insignificant, possibly because of endogenous sorting into the wallet type, that is, students who selected a bank-like wallet are already less worried about corporate surveillance to begin with. Adding the privacy preferences does not qualitatively change this finding (Column 4). Last, Columns (5) and (6) look at the students’ propensity to not link their Bitcoin wallet to a traditional bank account, making it more difficult for the government to tie their bitcoin transactions to their government-issued-currency ones. Consistent with the previous results, the encryption randomization – potentially because it gave the students a perception of protection from

initial interception – made it 3.8% less likely that the student would later try to escape surveillance from the government. When we interact the main effect with the students privacy preferences, results are consistent with privacy-sensitive students reacting less to the randomization, although estimates are very noisy.

Overall, while only suggestive because of the smaller sample size, taken together the results of this section highlight potential unexpected consequences of providing additional information on privacy protecting behavior. In our context, the discussion of how PGP technology can help consumers avoid initial interception - although irrelevant with respect to the privacy of future bitcoin transactions - seems to have increased disclosure by our participants towards the public, the intermediary and the government.

4 Conclusions

The privacy policy of both the US and OECD has focused on the idea that with enough transparency and enough choice consumers would make better privacy decisions. We explore consumers' attitude and revealed preferences towards digital privacy in the context of a large-scale experiment involving all MIT undergraduate students. We take advantage of the nature of the underlying technology (Bitcoin) to explore how this is moderated by preferences for privacy from a commercial firm, the government or the public.

Our results highlight a digital privacy paradox: Consumers say they care about privacy, but at multiple points in the process end up making choices that are inconsistent with their stated preferences.

The implications of our findings for policy are nuanced. Our finding that small incentives, costs or misdirection can lead people to safeguard their data less can have two interpretations. On the one hand it might lead policy makers to question the value of stated preferences for privacy when determining privacy policy. On the other hand, it might suggest the need for

more extensive privacy protections, from the standpoint that people need to be protected from their willingness to share data in exchange for relatively small monetary incentives.

Moreover, whenever privacy requires additional effort or comes at the cost of a less smooth user experience, participants are quick to abandon technology that would offer them greater protection. This suggests that privacy policy and regulation has to be careful about regulations that inadvertently lead consumers to be faced with additional effort or a less smooth experience in order to make a privacy-protective choice.

References

- Acquisti, A., C. Taylor, and L. Wagman (2016). The economics of privacy. *Journal of Economic Literature* 54(2), 442–492.
- Adjerid, I., A. Acquisti, L. Brandimarte, and G. Loewenstein (2013). Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, pp. 9:1–9:11. ACM.
- Allcott, H. and T. Rogers (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *The American Economic Review* 104(10), 3003–3037.
- Athey, S., I. Parashkevov, S. Sarukkai, and J. Xia (2016). Bitcoin pricing, adoption, and usage: Theory and evidence. Stanford University Graduate School of Business Research Paper.
- Barnes, S. B. (2006). A privacy paradox: Social networking in the United States. *First Monday* 11(9).
- Bertrand, M., D. Karlan, S. Mullainathan, E. Shafir, and J. Zinman (2010). What’s advertising content worth? Evidence from a consumer credit marketing field experiment. *The Quarterly Journal of Economics* 125(1), 263–306.
- Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2008). How are preferences revealed? *Journal of Public Economics* 92(8), 1787 – 1794.
- Catalini, C. and C. Tucker (2017). Seeding the S-curve? The role of early adopters in diffusion. MIT Sloan, Working Paper.
- Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *The American Economic Review* 99(4), 1145–1177.


- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature* 47(2), 315–372.
- DellaVigna, S., J. A. List, and U. Malmendier (2012). Testing for altruism and social pressure in charitable giving. *The Quarterly Journal of Economics* 127(1), 1–56.
- DellaVigna, S. and U. Malmendier (2006). Paying not to go to the gym. *The American Economic Review* 96(3), 694–719.
- Gneezy, U. and J. A. List (2006). Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments. *Econometrica* 74(5), 1365–1384.
- Gross, R. and A. Acquisti (2005). Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on Privacy in the Electronic Society, WPES '05*, New York, NY, USA, pp. 71–80. ACM: ACM.
- Harrison, G. W. and J. A. List (2004). Field experiments. *Journal of Economic Literature* 42(4), 1009–1055.
- Ho, D. E. and K. Imai (2008). Estimating causal effects of ballot order from a randomized natural experiment: The california alphabet lottery, 1978–2002. *Public Opinion Quarterly* 72(2), 216–240.
- Kim, J.-H. and L. Wagman (2015). Screening incentives and privacy protection in financial markets: A theoretical and empirical analysis. *The RAND Journal of Economics* 46(1), 1–22.
- Landry, C. E., A. Lange, J. A. List, M. K. Price, and N. G. Rupp (2010). Is a donor in hand better than two in the bush? Evidence from a natural field experiment. *The American Economic Review* 100(3), 958–983.

- Madrian, B. C. and D. F. Shea (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly Journal of Economics* 116(4), 1149–1187.
- Marotta-Wurgler, F. (2016). Self-regulation and competition in privacy policies. *The Journal of Legal Studies* 45(S2), S13–S39.
- McDonald, A. M. and L. F. Cranor (2009). The cost of reading privacy policies. *ISJLP* 4(3), 543–897.
- Miller, A. and C. Tucker (2011). Can healthcare information technology save babies? *Journal of Political Economy* 119(2), 289–324.
- Posner, R. A. (1981). The economics of privacy. *The American Economic Review* 71(2), 405–409.
- Reidenberg, J. R., J. Bhatia, T. D. Breaux, and T. B. Norton (2016). Ambiguity in privacy policies and the impact of regulation. *The Journal of Legal Studies* 45(S2), S163–S190.
- 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), 254–268.


5 Tables

Table 1: Descriptives


Variable	Mean	Std. Dev.	Min.	Max.	N
All Invalid (Ask)	0.06	0.238	0	1	3108
All Invalid (Ask + Incentive)	0.025	0.155	0	1	1543
Incentive Randomization	0.496	0.5	0	1	3108
Encryption Randomization	0.507	0.5	0	1	3108
Cash Out	0.394	0.489	0	1	3108
Bank-Like Wallet	0.713	0.452	0	1	3108
Year	2.457	1.11	1	4	3108
Male	0.551	0.497	0	1	3108
Top Coder	0.324	0.468	0	1	3108
Expected Price Decay	0.171	0.377	0	1	3108
Open Source Browser	0.125	0.331	0	1	3108
High Privacy from Peers	0.377	0.485	0	1	3108
High Privacy from Intermediary	0.548	0.498	0	1	3108
High Privacy from Government	0.424	0.494	0	1	3108
High Trust in Government	0.51	0.5	0	1	3108
High Trust in Startup	0.256	0.437	0	1	3108
High Trust in Retailer	0.434	0.496	0	1	3108
Selected Wallet Max. Priv. from Public	0.713	0.452	0	1	3108
Selected Wallet Max. Priv. from Intermediary	0.22	0.414	0	1	3108
Selected Wallet Max. Priv. from Government	0.091	0.288	0	1	3108
Wallet High Priv. Public Not Listed 1st	0.495	0.5	0	1	3108
Wallet High Priv. Intermediary Not Listed 1st	0.505	0.5	0	1	3108
Wallet High Priv. Government Not Listed 1st	0.753	0.431	0	1	3108
Escaping Surveillance from Public	0.018	0.131	0	1	1882
Escaping Surveillance from Intermediary	0.674	0.469	0	1	1410
Escaping Surveillance from Government	0.865	0.342	0	1	1410




(a) Ask versus Ask + Incentive



(b) By Stated Preferences for Privacy from the Public



(c) By Stated Preferences for Privacy from the Intermediary



(d) By Stated Preferences for Privacy from the Government

Figure 1: Effect of Small Incentives on Privacy (Comparison of Means)

Notes: The dependent variable in all figures is equal to one if the focal student only provided invalid entries when asked about the emails of her friends. 'Ask' refers to the non-incentivized condition, and 'Ask + Incentive' to the incentivized one.

Table 2: Effect of Small Incentives on Privacy

VARIABLES	(1) All Invalid	(2) All Invalid	(3) All Invalid	(4) All Invalid	(5) All Invalid	(6) All Invalid	(7) All Invalid
Ask + Incentive	-0.0285*** (0.0059)	-0.0268*** (0.0066)	-0.0224*** (0.0076)	-0.0249*** (0.0068)	-0.0245*** (0.0074)	-0.0327*** (0.0060)	-0.0332*** (0.0066)
Ask + Incentive × High Privacy Public		-0.0045 (0.0079)					
Ask + Incentive × High Privacy Intermediary			-0.0110 (0.0081)				
Ask + Incentive × High Privacy Government				-0.0085 (0.0078)			
Ask + Incentive × High Trust Government					-0.0080 (0.0079)		
Ask + Incentive × High Trust Startup						0.0153 (0.0100)	
Ask + Incentive × High Trust Retailer							0.0105 (0.0081)
Constant	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)
Observations	3,086	3,086	3,086	3,086	3,086	3,086	3,086
R-squared	0.005	0.006	0.006	0.006	0.006	0.006	0.006

Notes: The dependent variable in all columns is equal to one if the focal student only provided invalid entries when asked about the emails of her friends (either in the non-incentivized or in the incentivized condition). The sample only includes the 1,543 students (50% of participants, two observations per individual) that were randomly exposed both to the non-incentivized condition ('Ask'), and the incentivized one ('Ask + Incentive'). All columns show OLS regressions with robust standard errors clustered at the student level. Student-level variables capturing preferences for privacy and trust in different types of institutions are based on survey questions administered during registration. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Effect of Small Costs on Privacy

VARIABLES	(1) Maximized Privacy from the Public	(2) Maximized Privacy from the Public	(3) Maximized Privacy from the Public	(4) Maximized Privacy from the Intermediary	(5) Maximized Privacy from the Intermediary	(6) Maximized Privacy from the Intermediary	(7) Maximized Privacy from the Government	(8) Maximized Privacy from the Government	(9) Maximized Privacy from the Government
Best Wallet Not 1st	-0.1301*** (0.0161)	-0.1382*** (0.0204)	-0.1761*** (0.0240)	-0.1320*** (0.0147)	-0.1316*** (0.0221)	-0.1839*** (0.0225)	-0.0379*** (0.0130)	-0.0448*** (0.0167)	-0.0164 (0.0165)
High Stated Preference for Privacy		-0.0068 (0.0217)			-0.0156 (0.0232)			0.0066 (0.0240)	
Best Wallet Not 1st × High Stated Preference for Priv.		0.0216 (0.0332)			0.0000 (0.0296)			0.0157 (0.0266)	
Increased Transparency			0.0902*** (0.0210)			-0.1845*** (0.0225)			0.0644*** (0.0231)
Best Wallet Not 1st × Increased Transparency			0.0966*** (0.0318)			0.1082*** (0.0290)			-0.0391 (0.0257)
Constant	0.7777*** (0.0105)	0.7803*** (0.0133)	0.7313*** (0.0161)	0.2867*** (0.0115)	0.2951*** (0.0171)	0.3779*** (0.0174)	0.1199*** (0.0117)	0.1173*** (0.0151)	0.0859*** (0.0148)
Observations	3,108	3,108	3,108	3,108	3,108	3,108	3,108	3,108	3,108
R-squared	0.021	0.021	0.047	0.025	0.026	0.054	0.003	0.004	0.008

Notes: The dependent variable in Columns (1) to (3) is equal to one if the focal student selected a wallet that did not record all transactions on the public Bitcoin blockchain. In Columns (4) to (6), it is a binary indicator equal to one if the student selected an open-source wallet that does not allow the intermediary (if present) to access transaction data. In Columns (7) to (9), it is equal to one if the student selected Electrum, a wallet that is more difficult for the Government to track relative to the other featured ones. The key explanatory variable, 'Best Wallet Not 1st', is a binary indicator equal to one if none of the wallets that would maximize privacy on the focal dimension captured by the dependent variable is listed first on the wallet selection page. 'High Stated Preference for Privacy' is a dummy equal to one if the participant has high stated preference for privacy on the focal dimension (i.e. privacy from the public in Columns 1 to 3, from an intermediary in Columns 4 to 6, from the government in Columns 7 to 9). 'Increased Transparency' refers to the randomized condition (50% of the sample) in which the list of wallets was shown along with key privacy, security and convenience trade-offs. All columns use OLS regressions with robust standard errors. *** p<0.01, ** p<0.05, * p<0.1

Table 4: Effect of Small Talk on Privacy

VARIABLES	(1) Escape Surveillance from the Public (Coin Mixing)	(2) Escape Surveillance from the Public (Coin Mixing)	(3) Escaping Surveillance from the Intermediary (No Link To Phone/Address)	(4) Escaping Surveillance from the Intermediary (No Link To Phone/Address)	(5) Escaping Surveillance from the Government (No Link To Bank Account)	(6) Escaping Surveillance from the Government (No Link To Bank Account)
Encryption Randomization	-0.0105* (0.0061)	-0.0098 (0.0075)	-0.0361 (0.0250)	-0.0260 (0.0378)	-0.0331* (0.0182)	-0.0533** (0.0247)
High Stated Preference for Privacy		0.0038 (0.0103)		0.0308 (0.0356)		-0.0022 (0.0247)
Encryption Rand. \times High Stated Preference for Privacy		-0.0016 (0.0128)		-0.0169 (0.0504)		0.0475 (0.0363)
Constant	0.0229*** (0.0049)	0.0214*** (0.0061)	0.6921*** (0.0175)	0.6745*** (0.0272)	0.8820*** (0.0122)	0.8830*** (0.0162)
Observations	1,882	1,882	1,410	1,410	1,410	1,410
R-squared	0.002	0.002	0.001	0.002	0.002	0.005

Notes: The dependent variable in Columns (1) and (2) is equal to one if the focal students used a Bitcoin transaction mixing service to protect their privacy on the public Bitcoin blockchain. In Column (3) and (4), it is a binary indicator equal to one if they did not reveal their phone number or street address to the intermediary managing their digital wallet in the cloud. In Column (5) and (6), it is equal to one if the students did not connect their digital wallet to a traditional bank account, making it harder for the government to link their Bitcoin transactions to their government-issued-currency transactions. The 'Encryption Randomization', presented in Appendix Figure A1, was randomly shown to 50% of the sample. 'High Stated Preference for Privacy' is a dummy equal to one if the participant has high stated preference for privacy on the focal dimension (i.e. privacy from the public in Columns 1 and 2, from an intermediary in Columns 3 and 4, from the government in Columns 5 and 6). All columns use OLS regressions with robust standard errors. *** p<0.01, ** p<0.05, * p<0.1

6 Appendix

(a) Ask

(b) Ask + Incentive

Figure A-1: Privacy of Your Friends






Notes: The non-incentivized question ('Ask', Figure A-1a) was shown to the full sample. The incentivized one ('Ask + Incentive', Figure A-1b), was presented to a random, 50% of the sample.

Table A-1: Effect of Small Incentives on Privacy (Robustness)

VARIABLES	(1) All Invalid	(2) All Invalid	(3) All Invalid	(4) All Invalid	(5) All Invalid	(6) All Invalid	(7) All Invalid
Ask + Incentive	-0.0285*** (0.0059)	-0.0315*** (0.0083)	-0.0341*** (0.0071)	-0.0272*** (0.0064)	-0.0279*** (0.0064)	-0.0250** (0.0090)	-0.0292*** (0.0060)
Ask + Incentive × Year 2		0.0033 (0.0106)					
Ask + Incentive × Year 3		0.0022 (0.0106)					
Ask + Incentive × Year 4		0.0071 (0.0115)					
Ask + Incentive × Male			0.0100 (0.0078)				
Ask + Incentive × Top Coder				-0.0040 (0.0082)			
Ask + Incentive × Expected Price Decay					-0.0036 (0.0098)		
Ask + Incentive × Bank-Like Wallet						-0.0049 (0.0092)	
Ask + Incentive × Open Source Browser							0.0054 (0.0125)
Constant	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)	0.0531*** (0.0057)
Observations	3,086	3,086	3,086	3,086	3,086	3,086	3,086
R-squared	0.005	0.006	0.006	0.005	0.005	0.006	0.005

Notes: The dependent variable in all columns is equal to one if the focal student only provided invalid entries when asked about the emails of her friends (either in the non-incentivized or in the incentivized condition). The sample only includes the 1,543 students (50% of participants, two observations per individual) that were randomly exposed both to the non-incentivized condition ('Ask'), and the incentivized one ('Ask + Incentive'). All columns show OLS regressions with robust standard errors clustered at the student level. Student information (Year, Male) is provided by the MIT Office of the Provost, Institutional Research. Variables about the students' coding ability (Top Coder) and expectations about the future price of Bitcoin (Expected Price Decay) are collected through the registration survey. 'Bank-Like wallet' is equal to one if the students selected a digital wallet provided by a financial intermediary. 'Open Source Browser' is equal to one if the student signed up for the study using an open-source browser. Student-level variables capturing preferences for privacy and trust in different types of institutions are based on survey questions administered during registration. *** p<0.01, ** p<0.05, * p<0.1





Select your web browser(s)

 <p>Your online security is Firefox's top priority. Firefox is free, and made to help you get the most out of the web.</p>	 <p>Internet Explorer is the world's most widely used browser, designed by Microsoft with you in mind.</p>	 <p>The powerful and easy-to-use Web browser. Try the only browser with Opera Turbo technology, and speed up your Internet connection.</p>	 <p>Google Chrome. A fast new browser. Made for everyone.</p>	 <p>Safari for Windows from Apple, the world's most innovative browser.</p>
Install	Install	Install	Install	Install
Tell me more	Tell me more	Tell me more	Tell me more	Tell me more





[Further information, Terms of use and Privacy statement.](#)

Figure A-2: 'Browser Ballot' Screen

Notes: After the European Commission's ruling, Microsoft offered its users in Europe the choice between different, randomly sorted browsers. The 'ballot screen' presented each browser along with a short description and links to either install the browser or learn more about it. Source (March 2010): https://web.archive.org/web/20100323155508/http://www.browserchoice.eu/BrowserChoice/browserchoice_en.htm

Blockchain 	A hybrid web/self-managed wallet
Circle 	A web wallet service
Electrum 	A self-managed wallet
Coinbase 	A web wallet service
Other wallets	

(a) Example A

Coinbase 	A web wallet service
Electrum 	A self-managed wallet
Blockchain 	A hybrid web/self-managed wallet
Circle 	A web wallet service
Other wallets	

(b) Example B

Figure A-3: Wallet Order Randomization

Notes: The order of the four wallets was randomized. Each wallet was listed 1st for a random, 25% subset of our sample. The figure shows two of the possible combinations.



	How is my privacy protected?	How secure is my data?	How can I lose my bitcoin?	How can I access US dollars?	Can a government agency or the IRS seize my transaction data?
	<ul style="list-style-type: none"> + It is hard for external parties to identify you when you send out money - Has access to all your transactions 	<ul style="list-style-type: none"> + The company heavily invests in the security of their accounts - Your data is secure as long as any of the company security infrastructure is not breached, including through personal identity theft 	<ul style="list-style-type: none"> + If you are locked out of your wallet, you may be able to unlock it by proving your identity to the company - The company could freeze your account or go out of business 	<ul style="list-style-type: none"> + Easy conversion to and from US dollars (e.g., from bank account, credit card, debit card) 	- Yes
	<ul style="list-style-type: none"> + Only you have access to all your transactions - External parties could identify you when you send out money 	<ul style="list-style-type: none"> + Your data is secure as long as your passwords are secure 	<ul style="list-style-type: none"> + Nobody can freeze your account - If you are locked out of your wallet or lose access to your computer without a proper backup, it will be impossible for you to regain access 	<ul style="list-style-type: none"> - Only through a third-party service or individual 	+ No
	<ul style="list-style-type: none"> - Has access to all your transactions - External parties could identify you when you send out money 	<ul style="list-style-type: none"> + Your data is secure as long as your passwords are secure - Hackers could compromise your account through browser vulnerabilities 	<ul style="list-style-type: none"> + As long as you remember your password, you might be able to access a backup of your wallet - If you are locked out of your wallet, it may be impossible for you to regain access 	<ul style="list-style-type: none"> - Only through a third-party service or individual 	- Maybe
	<ul style="list-style-type: none"> + It is hard for external parties to identify you when you send out money - Has access to all your transactions 	<ul style="list-style-type: none"> + The company heavily invests in the security of their accounts - Your data is secure as long as any of the company security infrastructure is not breached, including through personal identity theft 	<ul style="list-style-type: none"> + If you are locked out of your wallet, you may be able to unlock it by proving your identity to the company + If the bitcoin is lost because of a breach at Circle, their deposit insurance may cover your loss - The company could freeze your account or go out of business 	<ul style="list-style-type: none"> + Easy conversion to and from US dollars (e.g., from bank account, credit card, debit card) 	- Yes
Other wallets					

Figure A-4: Wallet Choice under Increased Transparency

Notes: 50% of the sample ('Increased Transparency' condition) was randomly exposed to these columns which show key privacy, security and convenience trade-offs.

"PGP is a program that gives your electronic mail something that it otherwise doesn't have: Privacy. It does this by encrypting your mail so that nobody but the intended person can read it. When encrypted, the message looks like a meaningless jumble of random characters. PGP has proven itself quite capable of resisting even the most sophisticated forms of analysis aimed at reading the encrypted text.

PGP can also be used to apply a digital signature to a message without encrypting it. This is normally used in public postings where you don't want to hide what you are saying, but rather want to allow others to confirm that the message actually came from you. Once a digital signature is created, it is impossible for anyone to modify either the message or the signature without the modification being detected by PGP."

PGP makes sure that any communication between you and someone else can only be read by the sender and the receiver.

"End-to-end encryption creates a sort of digital tunnel between the senders and receivers of e-mails -- helping to keep the prying eyes of everyone from governments to Internet service providers and mail providers themselves from seeing the content of messages"

(source:

<http://www.washingtonpost.com/blogs/the-switch/wp/2014/08/07/yahoo-to-role-out-end-to-end-encryption-option-for-all-yahoo-mail-users-in-2015/>)

Although the technology has been available for a while, it is catching traction among those concerned about privacy and security. Both Yahoo! and Google have recently announced plans to integrate PGP into their email services. (For more information: http://en.wikipedia.org/wiki/Pretty_Good_Privacy)

Figure A-5: Encryption Randomization

Notes: All students saw the text at the top when we described PGP and encryption to them during the signup process. 50% were also randomly exposed to the bottom part (red box), which highlighted how PGP can be used to secure communication and avoid interception of the initial communication.