The Digital Privacy Paradox and Choice Architecture: Evidence from an Experiment in Fintech

Susan Athey, Christian Catalini, Alex Moehring, and Catherine Tucker*

August 30, 2025

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. We use data derived from a field experiment at MIT that distributed a new product, Bitcoin, to all undergraduates. There are two primary findings. First, small navigation costs have a tangible effect on how privacy-protective consumers' choices are, often in sharp contrast with individual stated preferences about privacy. Second, the introduction of irrelevant, but reassuring, information about privacy protection makes consumers less likely to avoid surveillance, regardless of their stated preferences toward privacy.

*Athey: Graduate School of Business, Stanford University, and NBER (athey@stanford.edu). Catalini: MIT Sloan School of Management, MIT (catalini@mit.edu). Moehring: Mitch Daniels School of Business, Purdue University. Tucker: MIT Sloan School of Management, MIT and NBER (cetucker@mit.edu).

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 if small design decisions by firms have substantial influence on consumer decisions to share personal information. In this paper, we present evidence of two potential distortions in the notice and choice process, relating to consumer decisions to share data and choose more or less privacy-protective technologies. In particular, we present evidence through a field experiment that simple design decisions that are ubiquitous in the digital economy have important impacts on consumer decisions to share personal information.

To do this, we use data from a digital currency experiment at the Massachusetts Institute of Technology where every undergraduate student was offered $100 in Bitcoin in the fall of 2014 (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 two digital privacy choices: Whether they choose a technology that maximizes 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 population is to small changes in navigation costs and information. This allows us to analyze how small frictions in the notice and choice process can have a large effect on privacy outcomes. In particular, we analyze how two seemingly innocuous design decisions can induce consumers to share more data than they otherwise would have: introducing small navigational costs and providing information.

The first 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 second treatment randomly exposed a subset of students to information about using encryption to protect their *initial* disclosure of information to us. We investigated whether the additional text made students more likely to protect their privacy from the public by obfuscating transactions on the Bitcoin public ledger, from the intermediary by not revealing additional identifying information to the wallet provider, or from the government by not linking their Bitcoin wallet to a traditional bank account subject to government oversight.

There are two main findings. First, small frictions introduced through navigation costs surrounding privacy choices can have large effects on technology adoption, even in the presence of information about the privacy consequences of those choices. Moreover, this impact of navigation costs persists among individuals who report caring a lot about privacy. Second, our information treatment on encryption – possibly by giving participants an illusion of protection – did not increase privacy-enhancing behavior, 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, consistent with the findings of Brandimarte et al. (2013). In all cases, privacy-decreasing decisions take place regardless of stated preferences for privacy.

Related Literature

This paper primarily contributes to three literatures. First, we contribute to the 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 that intentionally obfuscate their privacy notices (McDonald and Cranor, 2009), and that if presented with transparent information, consumers would make privacy choices that appear superior to those when they do not have transparent information (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). In contrast, our paper suggests that compliance with notice and consent may still not achieve the policy goal of protecting consumer privacy, given the many frictions and design decisions that a firm can exploit to induce users to share more data than they otherwise would have. Our paper demonstrates how these small frictions shift consumer behavior away from stated privacy preferences and how consent may differ from stated intent toward privacy.

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 second literature we contribute to is the literature studying 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, we define the privacy paradox as a disconnect between stated privacy beliefs and actual disclosure behavior. This definition is used in much of the privacy literature (Norberg et al., 2007; Acquisti, 2004; Gross and Acquisti, 2005; Barnes, 2006; John, 2015; Adjerid et al., 2018; Alashoor et al., 2023), though it is not universal (see Acquisti et al. (2020) for alternative definitions used).[1]

[1] For example, Solove (2021) argues that the privacy paradox is not a paradox at all and that stated preferences capture general privacy attitudes while revealed preferences reflect a specific scenario.

Much of this literature documents the privacy paradox in lab experiments (Norberg et al., 2007; Brandimarte et al., 2013; Adjerid et al., 2018; Alashoor et al., 2023) while a smaller literature documents the privacy paradox in field settings (Acquisti and Gross, 2006). This literature also explores factors that moderate the privacy paradox. For example Adjerid et al. (2013) shows, in the context of a lab experiment regarding a university social media site, that misdirection encouraged students to volunteer more sensitive academic data. Alashoor et al. (2023) find that emotional states and the extent of cognitive depletion moderate the privacy paradox. In particular, Alashoor et al. (2023) finds that individuals in states with low information processing ability or a positive mode display a larger disconnect between stated and revealed privacy preferences. Brandimarte et al. (2013) studies how perceived control over data access impacts current disclosure behavior. Our second study, where we randomize whether information on the encryption of participants' initial disclosure, tests this mechanism in a field setting.

The third literature we contribute to investigates privacy decisions through the lens of behavioral economics. This includes theoretical work investigating why privacy choices may differ from normative preferences (Acquisti, 2004). This literature also studies how small nudges can impact privacy choices such as framing, order effects, default settings, social norms, and anchoring (Johnson et al., 2002; Acquisti et al., 2012, 2013; Adjerid et al., 2019; Lin and Strulov-Shlain, 2025). This literature demonstrates how design and choice architecture can have large and important impacts on disclosure behavior.

We contribute to the existing literature in three ways. First, we contribute to the literature on the privacy paradox by demonstrating a disconnect between stated and revealed privacy preferences in a high-stakes field setting. Financial transactions can be highly sensitive and, due to this sensitivity, it is challenging to study privacy-related decisions in this setting. Using the unique field experiment that gives students $100 in Bitcoin we are able to study wallet choices that have important implications for the privacy of participants' transactions. In addition, we demonstrate the privacy paradox persists among highly sophisticated individuals.

Second, we contribute to the literature on behavioral economics and privacy decisions by investigating how recommendation systems and framing can influence privacy decisions. Much of the literature on behavioral economics and privacy decisions focuses on understanding mechanisms that explain disclosure behavior. We contribute by showing how choice-architecture or design can have large impacts on high-stakes privacy decisions.

Finally, we contribute to the policy debate surrounding notice and consent. In particular, we investigate how choice architecture, ranking in particular, can influence choices. We estimate a model of user privacy choices and simulate how choices differ between a ranking of wallets that optimizes for stated preferences and one that optimizes for disclosure.

Managerial and Policy Implications

This paper has important managerial and policy implications. For policymakers, 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 is slanted away from their actual normative preferences – as revealed by their stated privacy preferences in our surveys (Beshears et al., 2008). This could be a result of a number of behavioral biases that have been discussed in the literature, including present bias (Acquisti, 2004), information processing costs (Alashoor et al., 2023), incorrect beliefs about future control (Brandimarte et al., 2013), or the fact that consumers may not be able to accurately assess the value of their data (Acquisti et al., 2016) among other reasons (see John (2015) for a comprehensive review). We find that that subtle design choices, such as strategically ordering choices or providing additional but irrelevant information, can have large impacts on disclosure behavior. This brings into question the efficacy of the notice and choice framework.

The second policy stance our results can support is that there is a disconnect between stated and revealed preferences for privacy, but revealed preferences are actually closest to normative preferences. When expressing a preference for privacy is costless as it is in surveys, consumers are eager to express such a preference, but when faced with small navigational costs this taste for privacy quickly dissipates. This would suggest basing privacy protection on stated privacy preferences 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).

Importantly, the results of our study are positive and we do not make normative claims about the policy stances described above. Rather, we demonstrate in a high-stakes field setting how choice architecture can influence disclosure behavior. This is despite the fact that the experiment took place among an unusually well-informed and technologically aware population, suggesting that even highly sophisticated users often have a disconnect between stated privacy preferences and actual privacy choices. We believe this study represents a useful input for policymakers to make the normative judgment between the two interpretations.

For managers, our findings provide empirical evidence of how small design decisions can substantially alter consumer decisions to share personal information. There are clear analogs of the randomizations studied here in the digital economy. For example, firms often introduce small frictions in the choice process, such as the ordering and design of options on a web page or default choices. We find that by ordering offerings in a way that is designed to respect stated privacy preferences increases the probability a user makes a choice that aligns with their stated preferences by 7 percentage points relative to a nefarious design intended to maximize disclosure. This analysis has clear analogs to the tension between firms nudging users to choose options better for the users vs options better for the firm (Thaler, 2018). Second, firms often provide ample information about privacy, such as verbose privacy policies, which may be irrelevant to the actual privacy consequences of the choice. Our results suggest that these small design decisions can have large effects on consumer decisions to share personal information.

Empirical Setting and Data

In the fall of 2014, the MIT Bitcoin Club raised capital to give each of 4,494 MIT undergraduates $100 in Bitcoin. The aim of this experiment was not to study privacy. Instead, the objective was to jump-start the digital currency ecosystem on campus and expose their peers to the opportunities enabled by cryptocurrencies. By the end of our observation period (February 2016), the majority of participants (47.9%) were still holding on to their bitcoin, possibly because the cryptocurrency drastically increased in value (Catalini and Tucker, 2017).

As part of the sign-up 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. It is particularly useful that the purpose of the survey was unrelated to privacy. This allows us to study participant decisions to share personal information in a realistic field-experiment setting.

Our data covers the 3,108 undergraduates that signed up for a digital wallet. Participation by student class varied from 64% (fourth-years) to 79% (first-years). 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 with demographic information about the students provided by the Institutional Research section of the MIT Office of the Provost. Table 1 shows descriptive statistics for the sample.

To compare students' stated preferences for the privacy features of digital wallets with 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] We use the students' answers to divide the sample into high versus low taste for privacy from each of the three audiences. Students 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. 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 students' degree of trust in different institutions for financial services in the same way. The 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 students responded on a scale from 1 (not at all) to 5 (to a great extent). 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.

We analyze several randomizations that were included in the study given their implications for privacy. We first exploit the random order of wallets to understand how small navigation costs impact choices. This allows us to analyze the propensity of students to select a wallet that maximized privacy on different dimensions as a function of the wallet order on the page,

[2] The survey questions asked how important the privacy features of a digital wallet were on a scale from 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 in which the features were listed in was randomized.

Table 1: Descriptives

Variable	Mean	Std. Dev.	Min.	Max.	N
Encryption Randomization	0.507	0.500	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	1.110	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	0.510	0.500	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	0.287	0.452	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	0.500	0	1	3108
Wallet High Priv. Intermediary Not Listed 1st	0.505	0.500	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

which is exogenous by design. See Figures A-2a and A-2b for examples of this randomization. Cross-randomized with the order of wallets, half of the participants received additional detailed information about the privacy features of each wallet. Figure A-3 contains an example of the increased transparency choice screen.

A final randomization that we exploit is whether we include additional text on encryption. This 'Encryption Randomization' was randomly shown to 50% of the sample and is depicted in Figure A-4. 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 students took one of three actions to reduce the possibility of subsequent surveillance. The first outcome 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 blockchain. 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 fiat-currency transactions.

We test that the randomization was successful in two ways. First, Table A-1 presents balance tests showing the p-values for a joint test that different treatment groups have equal pre-treatment covariate means. The results of the tests presented in Table A-1 are consistent with randomization being successful. Second, we test randomization by testing that the expected share of individuals had a wallet that maximizes privacy from the public (50%), from an intermediary (50%), and from the government (75%) as the first in the random list of wallets. In addition, we can test if 50% of participants receive the PGP information treatment. We fail to reject the null in all cases that the expected share of individuals received each treatment, which is consistent with randomization being successful (Table 1).

Effect of Small Navigational Costs on Wallet Choice

During signup, students were shown four Bitcoin wallets, randomly ordered on the page. The randomized order in which the digital wallets were presented allows us to explore the effect of small frictions in a sign-up flow on 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 select and the data they end up disclosing. This process has a real world analogy: When searching for apps, consumers are confronted with a list of apps that may vary in the amount of data that they require the consumer to share despite performing similar functions.

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 regarding 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.[3]

Choices were strongly affected by the random ordering of wallets: When a bank-like

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

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 2 reports in regression format the effect of wallet order on technology choices. Columns (1) to (3) use an indicator equal to one if the 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 give an intermediary access to transaction data, and in Column (7) to (9) it is equal to one when students selected an open source wallet that is harder to track for the government. In each OLS regression the key explanatory variable, 'Privacy Max. Wallet Not 1st' is a binary indicator equal to one if the wallet listed first does not maximize privacy along the focal dimension. Specifically, the indicator 'Privacy Max. Wallet Not 1st' is equal to one when there are additional navigation costs to select the optimal wallet for the dimension of privacy captured by the dependent variable.

The results highlight how the navigational 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 percentage points 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, 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 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.

Table 2: Effect of Small Navigational Costs on Privacy

	Maximized Privacy from the Public			Maximized Privacy from the Intermediary			Maximized Privacy from the Government		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Privacy Max. Wallet Not 1st	-0.130***	-0.138***	-0.176***	-0.130***	-0.135***	-0.176***	-0.038***	-0.045***	-0.016
	(0.016)	(0.020)	(0.024)	(0.016)	(0.024)	(0.024)	(0.013)	(0.017)	(0.017)
High Stated Preference for Privacy		-0.007			-0.018			0.007	
		(0.022)			(0.024)			(0.024)	
Privacy Max. Wallet Not 1st ×		0.022			0.010			0.016	
High Stated Preference for Priv.		(0.033)			(0.032)			(0.027)	
Increased Transparency			0.090***			-0.187***			0.064***
			(0.021)			(0.024)			(0.023)
Privacy Max. Wallet Not 1st ×			0.097***			0.097***			-0.039
Increased Transparency			(0.032)			(0.032)			(0.026)
Intercept	0.778***	0.780***	0.731***	0.352***	0.362***	0.445***	0.120***	0.117***	0.086***
	(0.011)	(0.013)	(0.016)	(0.012)	(0.018)	(0.018)	(0.012)	(0.015)	(0.015)
Average of DV	0.713	0.713	0.713	0.287	0.287	0.287	0.091	0.091	0.091
Observations	3108	3108	3108	3108	3108	3108	3108	3108	3108
R-Squared	0.021	0.021	0.047	0.021	0.021	0.047	0.003	0.004	0.008

Notes: The dependent variable in Columns (1) to (3) is equal to one if the 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, 'Privacy Max. Wallet Not 1st', is a binary indicator equal to one if the wallet listed first on the wallet selection page does not maximize privacy along the focal dimension (i.e. Privacy Max. Wallet Not 1st equals 1 in Columns (1)-(3) if a the wallet that maximizes privacy from the public is not ranked first). '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. The row 'Average of DV' is the average dependent variable in each column. All columns use OLS regressions with robust standard errors. *** p<0.01, ** p<0.05, * p<0.1

Figure A-3 shows a screenshot of the randomized treatment we exploit in this analysis: Whereas 50% of the sample only saw names, logos, and short descriptions of the four randomly sorted wallets (see Figures A-2a and A-2b 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-3), 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 2 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 privacy-maximizing wallet was listed first and no additional information was provided, as before, a lower ranking substantially decreases the likelihood that the participant will maximize privacy along the focal dimension. If we focus on point estimates, these results suggest that, while additional transparency dilutes the effect of navigation costs, it does not 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 percentage points), but this time maps to a larger relative effect, as the probability of making this type of choice is only 28.7%. When a wallet that optimizes privacy from an intermediary is not listed first, we observe a 36.9% 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 privacy maximizing 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 percentage point 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, Table 2 supports the idea that the order of wallets on the sign-up page, by introducing minor costs for the wallets not ranked first, had a large effect on the technology adopted by students across three different privacy dimensions. 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 tastes 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.

Relative Importance of Wallet Order, Preferences, and Information

We now estimate a series of discrete-choice models of participant wallet choices. While the linear-probability models presented above demonstrate that wallet order impacts choice over wallet, the discrete-choice models will allow us to decompose the relative importance of participant preferences over wallet features and wallet order.

Assume participant i receives utility from selecting wallet j that depends on the position, or rank, of the wallet in the list ($r_j \in \{0, 1, 2, 3\}$), a Likert-scale of participant i's stated preference for features contained by wallet j ($s_{ij} \in \{1, , 2, 3, 4, 5\}$), and whether participant i was randomized into the additional transparency condition ($t_i \in \{0, 1\}$):

$$U_{ij} = u_j\left(r_{ij}, s_{ij}, t_i\right) + \varepsilon_{ij}. \tag{1}$$

Participants choose the wallet that maximizes utility and there is no outside option as all participants choose a wallet. In particular, a user chooses wallet j if $U_{ij} \geq U_{ij'}$ for all $j' \neq j$. We take this model to the data by imposing a linear functional form for $u_j(\cdot)$ with an indicator if wallet j is a bank-like wallet and interactions between the transparency condition and r_{ij}, s_{ij}, and the bank-like indicator. We further assume the preference shock, ε_{ij}, is IID type-1 extreme value and estimate the conditional logit model using maximum likelihood (Train, 2009).

While wallet position likely impacts choice by reducing the probability that an individual considers the model, a fully-specified search model is beyond the scope of this exercise. Rather, we include wallet rank directly in the utility function as a reduced-form approach to capturing the impact position has on choice.

Table 3 presents the preference estimates from this model. Column (1) estimates a model that only considers whether a wallet is bank-like or not. As expected from the raw choice probabilities, consumers strongly prefer bank-like wallets. Column (2) adds the impact of wallet position and it is clear that consumers have a strong preference for wallets ranked earlier in the list. On average, swapping the position of the first and last wallet's position decreases the predicted choice probability of the wallet originally positioned first by 12.7 percentage points.

Column (3) further adds s_{ij} to the model, which captures whether a wallet contains features for which the user stated they have a strong preference. Stated preferences are predictive of choice though we are unable to reject that the preference for wallets with a high stated preferences is 0. On average, changing the stated preference for a wallet to its maximum value (5) increases the predicted choice probability by only 2.4 percentage points. Finally, Column (4) allows for different preferences between the group randomized into receiving additional transparency and those without. We see that increasing transparency substantially reduces the impact of wallet position on choice probabilities, nearly cutting the preference for position in half. Moreover, increased transparency induces more individuals to choose bank-like wallets.

These findings can be rationalized by a model where consumers have uncertainty over the privacy features of each wallet and a preference for convenience (i.e. wallets positioned higher up the page). Without additional information, consumers are more likely on average to choose wallets aligning with their stated preferences. However, the convenience factor is also important for these consumers and leads them to choose sub-optimal wallets from a stated preference perspective. Providing additional information to consumers allows them to better assess each wallet and reduces the impact of wallet position on choice.

In addition to allowing us to decompose the relative importance of stated privacy preferences and wallet order on choice, this model also allows us to predict choices of counterfactual wallet orderings. We consider two counterfactual orderings. First, we consider a firm whose goal is to respect individuals' stated preferences while maintaining user autonomy over choices. This firm ranks wallets to maximize the probability that the chosen wallet is the wallet that maximizes privacy along the user's most preferred dimension. Second, we consider a firm who wants to minimize the probability the user chooses a wallet that maximizes privacy along the user's most preferred dimension. We can think of the latter counterfactual as a nefarious firm who wants to maximize disclosure among the user's most sensitive dimension. We summarize the results of this counterfactual by comparing the probability that a user chooses their most preferred wallet based on their stated preferences. The random ordering of wallets users encounter in the experiment results in 64.5% of users choosing a wallet that maximizes privacy among the dimension they report caring the most about. This number rises to 68.1% for a firm who orders wallets to respect stated preferences. For the nefarious firm, however, this number falls to 61.3%. This 7 percentage point difference highlights how sensitive high-stakes privacy decisions are to one of the many subtle design decisions available to the firm.

Table 3: Decomposing Impact of Stated Preferences and Wallet Order on Choice

	(1)	(2)	(3)	(4)
Bank-Like Wallet	0.912^{***}	0.918^{***}	0.939^{***}	0.616^{***}
	(0.040)	(0.040)	(0.043)	(0.057)
Wallet Order		-0.181^{***}	-0.152^{***}	-0.193^{***}
		(0.017)	(0.045)	(0.048)
Stated Preference			0.061	0.045
			(0.039)	(0.049)
Stated Preference ×			-0.010	-0.010
Wallet Order			(0.015)	(0.015)
Bank-Like Wallet ×				0.694^{***}
Increased Transparency				(0.087)
Wallet Order ×				0.083^{**}
Increased Transparency				(0.034)
Stated Preference ×				0.031
Increased Transparency				(0.068)
Observations	12432	12432	12432	12432
R-Squared				

Notes: Maximum-likelihood estimates of Equation 1 for various utility specifications. The log-likelihood is given by: $\mathcal{L}(\beta) = \sum_i \sum_j y_{ij} \log P_{ij}(\beta)$ where y_{ij} is an indicator equal to 1 if user i selected wallet j, P_{ij} is the logit choice probability, and β is the preference vector. Wallet Order is the position of the wallet in the list (0-3). Stated Preference is the participant's stated preference (1-5) of the privacy features of a given wallet. Bank-Like Wallet is an indicator equal to one for bank-like wallets. Increased Transparency is an indicator equal to 1 if individual i received the additional transparency randomization. All standard errors are clustered at the individual-level. There are 4 × Number of participant observations as each individual chooses from four wallet options.

The Role of Unrelated Information in Nudging Choices

In this section, we study the impact of an information treatment that explicitly focused on privacy-protecting behavior. Under the 'Encryption Randomization' condition (Figure A-4), when shown details about PGP technology, 50% of the sample received 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 the previous analysis.

Although PGP encryption technology is widely used in contexts where security and privacy are paramount, in our setting the technology did not provide students with any additional protection of future Bitcoin transactions. These would be still exposed to intermediaries, governments, or the public depending on the digital wallet selected by 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)-(3), 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 transactions 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 how to use it, the baseline is very low. Only 2.3% of students used a mixing service when not exposed to the 'Encryption Randomization.' The percentage goes further down among the treated group, which corresponds to a 1 percentage point reduction in use and a 45% relative decrease. Results are noisy and we may be unable to capture some of the most sophisticated methods of transaction mixing or the use of less popular mixing services. Although the 1 percentage point reduction in use may seem small, it is important to remember that overall activity with Bitcoin is also low, at 13.1%. Column (2) interacts the encryption randomization with a self-reported measure of coding ability. We find that the decline in the use of mixing services in response to the encryption randomization is concentrated among Top Coders. While we are underpowered to make strong claims here, we believe this is suggestive that sophisticated individuals may be more susceptible to the control paradox (Brandimarte et al., 2013), where perceived control over future data access impacts current disclosure behavior. In Column (3), we interact the treatment with students' stated privacy preferences and results are qualitatively unchanged (the main effect for the randomization is insignificant because of the smaller sample size, but comparable in size and standard error to the effect in Column 1).

Columns (4) to (9) of Table 4 further limit the sample to 'Bank-Like' wallets, as for these wallets we can observe students' decisions to disclose additional information that may make it easier for the intermediary or the government to track. Table A-2 shows versions of Columns (1)-(3) on this same sample and the results are consistent.

The dependent variable in Columns (4)-(6) 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 encryption randomization in Column (4) is negative but small, noisy, and statistically 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 privacy preferences does not qualitatively change this finding (Column 6). However, we again see a statistically insignificant but economically large treatment effect among sophisticated users (Top Coders) in Column (5).

Finally, Columns (7)-(9) 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 transaction. Consistent with previous results, the encryption randomization – potentially because it gave the students a perception of protection from initial interception – made it 3.8% less likely (3.3 percentage points) that students would later try to escape surveillance from the government. This is again concentrated among sophisticated users (Column 8), though the heterogeneous effect is not statistically significant. When we interact the main effect with privacy preferences, results are consistent with privacy-sensitive students reacting less to the randomization, although estimates are noisy.

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.

Conclusion

The privacy policies of many industrial nations have focused on the idea that with enough transparency and choice consumers would make better privacy decisions. We explore consumers' attitudes and revealed preferences towards digital privacy in the context of a large-scale field experiment involving all MIT undergraduate students. We also 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 navigation costs or misdirection can lead people to safeguard their data less can have two interpretations. On the one hand, it might lead policymakers 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 small monetary incentives.

Moreover, when privacy requires additional effort or degrades user experience, participants are quick to abandon technology that offers greater protection. This suggests that privacy policy and regulation must be careful about regulations that inadvertently require additional effort or a less smooth experience in order to make a privacy-protective choice.

There are limitations to our study. First, the study was conducted within MIT, which is an environment where people are more technologically and data-aware than many settings. Though this implies we abstract from some of the 'knowledge' constraints that other privacy studies face, it also means we do not study the actions of unsophisticated consumers. Second, the behavioral insights that we document such as the role of information, order effects, and the inconsistency between stated and revealed preferences are not new behavioral phenomenon but instead have been widely documented in other settings. Finally, we acknowledge that Bitcoin was a relatively new technology during the time of the study. It is possible that as individuals become more familiar with the technology, the barrier to making choices that preserve their privacy is reduced, so our results should be interpreted as being about educated and technologically savvy people when faced with an unfamiliar technology. Notwithstanding these limitations, we believe our study is a useful exploration of the privacy paradox in the field.

Table 4: Effect of Extraneous Information on Privacy Choices

	Remove Visibility from Public			Remove Visibility from Intermediary			Remove Visibility from Government		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Encryption Randomization	-0.010*	-0.006	-0.010	-0.036	-0.008	-0.026	-0.033*	-0.015	-0.053**
	(0.006)	(0.007)	(0.008)	(0.025)	(0.030)	(0.038)	(0.018)	(0.021)	(0.025)
Top Coder		0.005			0.056			-0.010	
		(0.011)			(0.037)			(0.026)	
Encryption Randomization × Top Coder		-0.013			-0.087			-0.062	
		(0.013)			(0.054)			(0.041)	
High Stated Preference for Privacy			0.004			0.031			-0.002
			(0.010)			(0.036)			(0.025)
Encryption Randomization × High Stated Preference for Priv.			-0.002			-0.017			0.047
			(0.013)			(0.050)			(0.036)
Intercept	0.023***	0.021***	0.021***	0.692***	0.673***	0.674***	0.882***	0.885***	0.883***
	(0.005)	(0.006)	(0.006)	(0.018)	(0.022)	(0.027)	(0.012)	(0.015)	(0.016)
Average of DV	0.018	0.018	0.018	0.674	0.674	0.674	0.865	0.865	0.865
Observations	1882	1882	1882	1410	1410	1410	1410	1410	1410
R-Squared	0.002	0.002	0.002	0.001	0.004	0.002	0.002	0.007	0.005

Notes: The dependent variable in Columns (1)-(3) is equal to one if the student used a Bitcoin transaction mixing service to protect their privacy on the public Bitcoin blockchain. In Columns (4)-(6), 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 Columns (7)-(9), 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). 'Top Coder' is a dummy equal to one if the participant has strong self-assessed programming skills. All columns restrict to individuals who did not immediately cash-out their Bitcoin and Columns (4)-(9) further restrict to individuals who selected a bank-like wallet. The row 'Average of DV' is the average dependent variable in each column. All columns use OLS regressions with robust standard errors. *** p<0.01, ** p<0.05, * p<0.1

References

Acquisti, A. (2004). Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 21–29.

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

Acquisti, A. and Gross, R. (2006). Imagined communities: Awareness, information sharing, and privacy on the facebook. In *International workshop on privacy enhancing technologies*, pages 36–58. Springer.

Acquisti, A., John, L. K., and Loewenstein, G. (2012). The impact of relative standards on the propensity to disclose. *Journal of Marketing Research*, 49(2):160–174.

Acquisti, A., John, L. K., and Loewenstein, G. (2013). What is privacy worth? *The Journal of Legal Studies*, 42(2):249–274.

Acquisti, A., Taylor, C., and Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2):442–492.

Adjerid, I., Acquisti, A., Brandimarte, L., and Loewenstein, G. (2013). Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, pages 9:1–9:11. ACM.

Adjerid, I., Acquisti, A., and Loewenstein, G. (2019). Choice architecture, framing, and cascaded privacy choices. *Management science*, 65(5):2267–2290.

Adjerid, I., Peer, E., and Acquisti, A. (2018). Beyond the privacy paradox: Objective versus relative risk in privacy decision making. *MIS quarterly*, 42(2):465–488.

Alashoor, T., Keil, M., Smith, H. J., and McConnell, A. R. (2023). Too tired and in too good of a mood to worry about privacy: explaining the privacy paradox through the lens of effort level in information processing. *Information Systems Research*, 34(4):1415–1436.

Athey, S., Parashkevov, I., Sarukkai, S., and Xia, J. (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).

Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2008). How are preferences revealed? *Journal of Public Economics*, 92(8):1787 – 1794.

Brandimarte, L., Acquisti, A., and Loewenstein, G. (2013). Misplaced confidences: Privacy and the control paradox. *Social psychological and personality science*, 4(3):340–347.

Catalini, C. and Tucker, C. (2017). When early adopters don't adopt. *Science*, 357(6347):135–136.

Gross, R. and Acquisti, A. (2005). Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on Privacy in the Electronic Society*, WPES '05, pages 71–80, New York, NY, USA. ACM, ACM.

John, L. K. (2015). The consumer psychology of online privacy: Insights and opportunities from behavioral decision theory. *The Cambridge Handbook of Consumer Psychology*.

Johnson, E. J., Bellman, S., and Lohse, G. L. (2002). Defaults, framing and privacy: Why opting in-opting out. *Marketing letters*, 13:5–15.

Kim, J.-H. and Wagman, L. (2015). Screening incentives and privacy protection in financial markets: A theoretical and empirical analysis. *The RAND Journal of Economics*, 46(1):1–22.

Lin, T. and Strulov-Shlain, A. (2025). Choice architecture, privacy valuations, and selection bias in consumer data. *Marketing Science*.

Marotta-Wurgler, F. (2016). Self-regulation and competition in privacy policies. *The Journal* *of Legal Studies*, 45(S2):S13–S39.

McDonald, A. M. and Cranor, L. F. (2009). The cost of reading privacy policies. *ISJLP*, 4(3):543–897.

Miller, A. and Tucker, C. (2011). Can healthcare information technology save babies? *Journal of Political Economy*, 119(2):289–324.

Norberg, P. A., Horne, D. R., and Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of consumer affairs*, 41(1):100–126.

Posner, R. A. (1981). The economics of privacy. *The American Economic Review*, 71(2):405–409.

Reidenberg, J. R., Bhatia, J., Breaux, T. D., and Norton, T. B. (2016). Ambiguity in privacy policies and the impact of regulation. *The Journal of Legal Studies*, 45(S2):S163–S190.

Solove, D. J. (2021). The myth of the privacy paradox. *Geo. Wash. L. Rev.*, 89:1.

Thaler, R. H. (2018). Nudge, not sludge.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

Tsai, J. Y., Egelman, S., Cranor, L., and Acquisti, A. (2011). The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, 22(2):254–268.

For Online Publication: Appendix

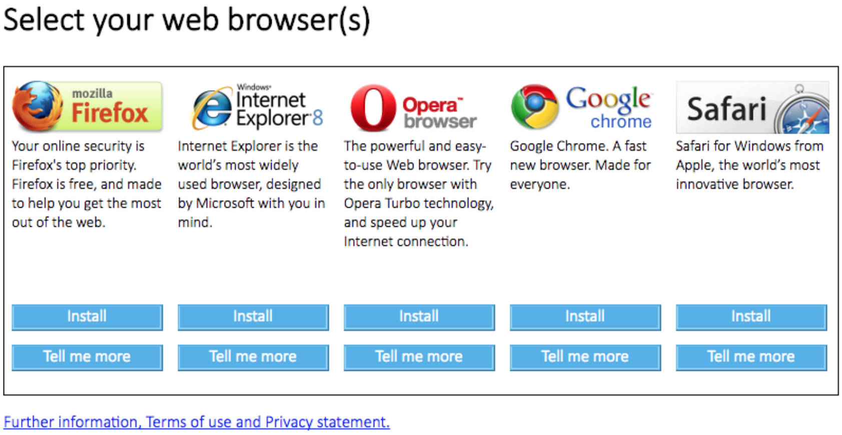

Select your web browser(s)

mozilla Firefox	Windows Internet Explorer 8	Opera browser	Google chrome	Safari
Your online security is Firefox's top priority. Firefox is free, and made to help you get the most out of the web.	Internet Explorer is the world's most widely used browser, designed by Microsoft with you in mind.	The powerful and easy-to-use Web browser. Try the only browser with Opera Turbo technology, and speed up your Internet connection.	Google Chrome. A fast new browser. Made for everyone.	Safari for Windows from Apple, the world's most innovative browser.
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-1: '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

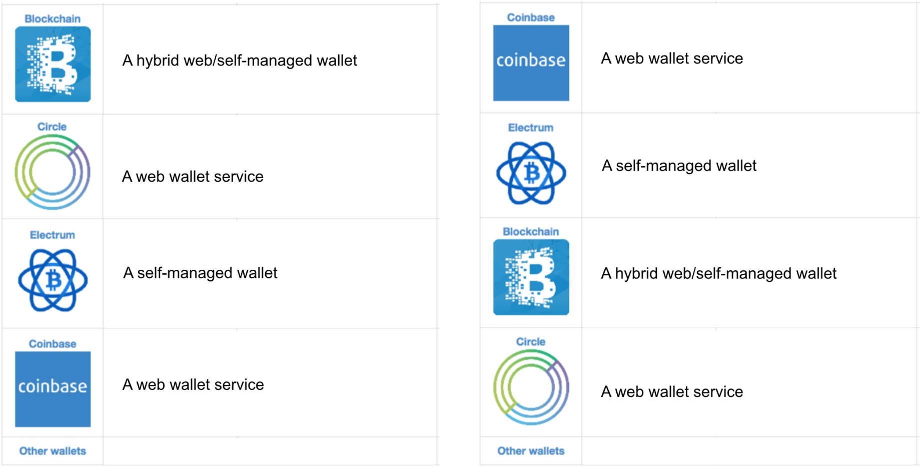

Blockchain	A hybrid web/self-managed wallet
Circle	A web wallet service
Electrum	A self-managed wallet
Coinbase	A web wallet service
Other wallets	

(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-2: 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.

Table A-1: Randomization Check - Covariate Balance

Covariate	Small cost	Small talk
Year	0.654	0.128
Male	0.929	0.446
Top Coder	0.400	0.063
Expected Price Decay	0.175	0.165
Open Source Browser	0.348	0.716
High Privacy from Peers	0.830	0.533
High Privacy from Intermediary	0.382	0.857
High Privacy from Government	0.134	0.833
High Trust in Government	0.022	0.960
High Trust in Startup	0.242	0.066
High Trust in Retailer	0.969	0.333
Escaping Surveillance from Public	0.366	0.086
Escaping Surveillance from Intermediary	0.237	0.148
Escaping Surveillance from Government	0.403	0.069

Notes: This table reports the p-value for a test of the null hypothesis that the mean of each pre-treatment covariate among the groups receiving treatment and those that do not are equal. For the small navigational cost condition, this is a joint test that the mean covariates are equal among people who received the wallet maximizing privacy from the public, intermediary, and the government.

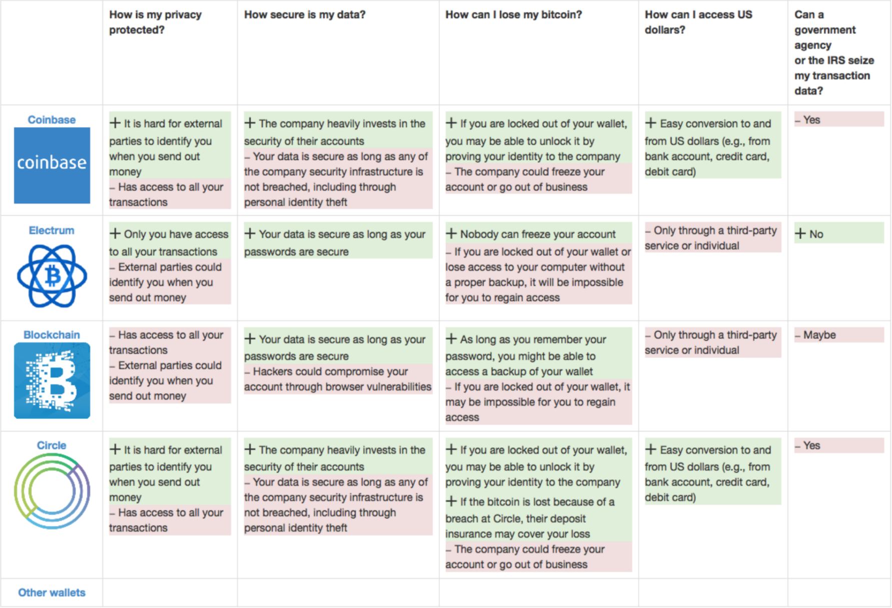

	How is my privacy protected?	How secure is my data?	How can I lose my bitcoin?	How can I access US dollars?	Can a government agency or the IRS seize my transaction data?
Coinbase	+ It is hard for external parties to identify you when you send out money – Has access to all your transactions	+ The company 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	+ 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	+ Easy conversion to and from US dollars (e.g., from bank account, credit card, debit card)	– Yes
Electrum	+ Only you have access to all your transactions – External parties could identify you when you send out money	+ Your data is secure as long as your passwords are secure	+ 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	– Only through a third-party service or individual	+ No
Blockchain	– Has access to all your transactions – External parties could identify you when you send out money	+ Your data is secure as long as your passwords are secure – Hackers could compromise your account through browser vulnerabilities	+ 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	– Only through a third-party service or individual	– Maybe
Circle	+ It is hard for external parties to identify you when you send out money – Has access to all your transactions	+ 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	+ 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	+ Easy conversion to and from US dollars (e.g., from bank account, credit card, debit card)	– Yes
Other wallets					

Figure A-3: 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.

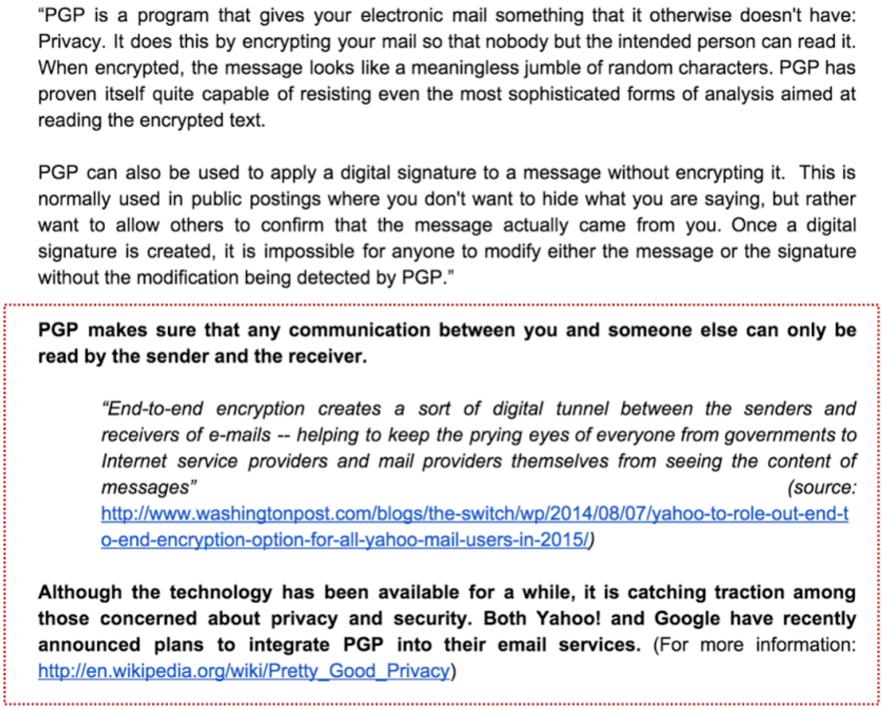

"PGP is a program that gives your electronic mail something that it otherwise doesn't have: Privacy. It does this by encrypting your mail so that nobody but the intended person can read it. When encrypted, the message looks like a meaningless jumble of random characters. PGP has proven itself quite capable 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-4: 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.

Table A-2: Effect of Extraneous Information on Privacy Choices - Robustness to Sample Selection

	Remove Visibility from Public		
	(1)	(2)	(3)
Encryption Randomization	-0.008	-0.002	-0.009
	(0.007)	(0.008)	(0.009)
Top Coder		0.013	
		(0.013)	
Encryption Randomization × Top Coder		-0.019	
		(0.015)	
High Stated Preference for Privacy			-0.000
			(0.011)
Encryption Randomization × High Stated Preference for Priv.			0.003
			(0.015)
Intercept	0.022***	0.017***	0.022***
	(0.006)	(0.006)	(0.007)
Average of DV	0.018	0.018	0.018
Observations	1410	1410	1410
R-Squared	0.001	0.002	0.001

Notes: This table presents treatment effects of the encryption randomization on an indicator equal to one if the student used a Bitcoin transaction mixing service to protect their privacy on the public Bitcoin blockchain among individuals who used a bank-like wallet and did not immediately cash-out their Bitcoin. The row 'Average of DV' is the average dependent variable in each column. All columns use OLS regressions with robust standard errors. *** p<0.01, ** p<0.05, * p<0.1