

AN ONLINE EXPERIMENT ON CONSUMERS' SUSCEPTIBILITY TO FALL FOR POST-TRANSACTION MARKETING SCAMS

Complete Research

Nochenson, Alan, The Pennsylvania State University, College of Information Sciences and Technology, University Park, USA, anochenson@psu.edu

Grossklags, Jens, The Pennsylvania State University, College of Information Sciences and Technology, University Park, USA, jensg@ist.psu.edu

Abstract

Post-transaction marketing offers are often designed to trick consumers into purchasing products they would not want. To increase the frequency of transactions, retailers use strategies such as subscribing consumers by default to offers, camouflaging post-transaction offers as part of primary transactions, and being unclear about the terms of the offer. Further, sharing agreements for personal consumer data and payment credentials between first and third-party retailers violate consumers' expectations about privacy and business conduct. Some of these tactics have been mitigated in the United States with the Restore Online Shoppers' Confidence Act (ROSCA). In this paper, we report the results from an online post-transaction marketing experiment. We involve over 500 consumers in a purchasing episode (i.e., a functional mock-up music store) followed by a post-transaction marketing offer with low value to consumers. In the experiment, we systematically vary important design characteristics on the offer page, and collect additional data through a post-experimental questionnaire. We investigate which factors are most predictive of acceptance of the low-value post-transaction offer. We find that ROSCA's interventions are a step in the right direction and should be considered by European regulators, but do not go far enough.

Keywords: Post-Transaction Offer Marketing, Online Experiment, Computer-Human Interaction, Protection of Personal Information.

1 Introduction

Post-transaction marketing offers are presented to consumers after they have successfully concluded a primary purchase. The objective is to encourage an additional transaction with the consumer and not merely to advertise a different product. In principle, such practices can be part of the normal and broadly acceptable interactions between businesses and consumers, in particular, if the offered products and services are sufficiently related to the primary purchase and will likely appeal to consumers.

Unfortunately, many examples of post-transaction marketing have to be considered potentially deceptive and misleading according to the definition given by the Federal Trade Commission Policy Statement on Deception (FTC, 1984). In fact, until 2008, the FTC brought cases against over 300 companies and individuals (Ballaré and Von Bergen, 2008) for practices that are “likely to mislead the...reasonable consumer” (FTC, 1984). To investigate the extent of such practices, the U.S. Senate compiled a report that detailed how nearly 4 million customers were enrolled in programs that had been sold as potentially deceptive post-transaction offers in 2009, and over 35 million consumers were affected in the previous 10-year period. Even more worrisome, 99% of the affected consumers (that were surveyed for the report) had never used the product, were typically unaware of having purchased the product, and were accordingly not satisfied with their purchase (U.S. Senate, 2009).

In most investigated cases, post-transaction marketing offers were provided by third-party entities that formed a profit-sharing relationship with well-known businesses. Under the terms of these business relationships, a consumer would be automatically redirected to a third-party site after the conclusion of the primary purchase. However, consumers likely did not realize that such a redirection had taken place and believed that the post-transaction marketing page was merely a continuation of the check-out process (Edelman, 2009).

In addition, the investigated business relationships involved different degrees of sharing of consumer data and payment credentials between the primary merchant (who collected the data from the consumer) and the third party (which had in most cases no previous relationship with the affected consumer). The information was used to pre-populate the post-transaction offer with personal information and payment details to lower the bar for agreement to the offer. However, these so-called data-pass arrangements violated consumers’ expectations about data protection and data privacy, and about the typical treatment of consumer information during electronic commerce transactions (Acquisti and Grossklags, 2005; Turow et al., 2007). Further, data-pass also contributed to the mistaken assumption by consumers that the primary transaction was still ongoing (Edelman, 2009).

In response to these findings, U.S. Congress passed the Restore Online Shoppers’ Confidence Act (ROSCA, 2010). ROSCA prohibits post-transaction offers unless they “clearly and conspicuously disclose all...terms of the transaction,” which include a description of what the offer is, its price, and the involved parties. To charge consumers, merchants are compelled to ensure consumers “perform an additional affirmative action such as clicking on a confirmation button or checking a box that indicates...consent to be charged.” Accordingly, ROSCA outlaws the most aggressive forms of post-transaction marketing in which consumers were presumed to have agreed to the offer by default (and in some cases did not even have an option to cancel the order on the post-transaction marketing site). However, it remains unclear whether the restrictions compelled by ROSCA significantly lower the likelihood that consumers will fall for offers that they do not want and do not need.

To the best of our knowledge, we report the first comprehensive experimental study that investigates consumers’ response to post-transaction marketing offers. We developed a plausible primary online shopping scenario and a related post-transaction marketing offer that we anticipated to have very low value to an average consumer. Further, decisions in the experiment had actual payoff consequences for the participants.

In our study, we systematically vary three important design dimensions that are relevant in the context of post-transactions marketing. First, we investigate the impact of different default regimes for agreement to the offer. We consider opt-in and opt-out mechanisms, i.e., in the latter case the participant has to take affirmative action when she does not want to accept the offer. Second, we study different data protection scenarios. That is, we compare experimentally an arrangement with data-pass and a scenario in which the participant has to enter her information on the post-transaction page herself. Third, we investigate how the visibility of data-pass arrangements impacts consumer behaviour, i.e., we also study scenarios in which data is shared with the third party; however, the shared information is not made visible on the post-transaction marketing page. The results of our study allow for a discussion of the positive impact and the limitations of ROSCA and similar legislation, and shed light at the complex task to protect consumers from aggressive online marketing practices which violate data protection expectations in unanticipated ways.

In Section 2, we provide additional background information about post-transaction marketing and review related work. In Section 3, we present our experimental methodology. Section 4 contains the analysis of the experimental data. In Section 5, we discuss the results and offer concluding remarks.

2 Background

2.1 Post-Transaction Marketing in Practice

Post-transaction marketing sales are driven by a mixture of factors that may include “data pass” arrangements (i.e., payment credentials are automatically transferred between partners and are not re-entered by the consumer), “free-to-pay” conversions (i.e., free trial offers that without additional consumer interactions result in payments), “negative options” (i.e., perpetual membership that causes periodic payments without further notice), and “diffusion” (e.g., obstructions on the ability of consumers to identify the relevant parties). Users severely struggle with the powerful combination of these factors (Cox, 2010; U.S. Senate, 2009).

Senate Committee hearings in 2009 showed that research conducted in other contexts exists to highlight aspects of the problem, for example, from a readability or high level consumer decision-making perspective (U.S. Senate, 2009). Further, prominent researchers in information design have recognized many related problematic issues such as the intentional masking of substantive content (e.g., Tufte, 1997). However, we are unaware of research that systematically studies these issues within the specific context of post-transaction marketing or sufficiently related areas to formulate appropriate public policy responses, to design user support tools, and to educate consumers to implement their genuine desires in the marketplace. Our assessment is supported by comprehensive review efforts in marketing and decision making (Boush et al., 2009; Milkman et al., 2009).

In our online experiment, we research several (but not all) characteristics of post-transaction marketing. In particular, we systematically vary the process and visibility of data pass arrangements, and the process of gaining user consent (e.g., opt-in/opt-out). We keep a consistent level of diffusion across our experimental treatments and use a design of the post-transaction marketing that is derived from legal cases (e.g. the Intelius case; see Johnson 2010). The experiment is situated in an actual purchasing situation with monetary consequences to the participants. Altogether, the experimental setup is conducive to provide actionable insights for the public policy process and to yield input for the development of best-practice marketing activities.

2.2 Online Notice and Consent

Our study is influenced by online notice and consent research. Most post-transaction marketing schemes typically disclose the most relevant terms, though the ways that they do so are usually unexpected and difficult to comprehend for an average consumer. Similarly, other types of online

notices and most privacy policies present complex legal information that most consumers find impenetrable, or at the very least user-unfriendly.

Research shows that complexity of privacy notices hampers users' ability to understand such agreements. For example, Jensen and Potts (2004) studied a sample of 64 privacy policies from high traffic and health care websites. They found that policies' format, location on the website and legal content severely limit users' ability to make informed decisions. Good et al. (2007) conducted an experiment in which they presented a carefully developed short-notice before or after an installation choice. The availability of this more concise information reduced the occurrence of installations of problematic programs (compared to a control treatment of a standard notice). Somewhat surprisingly, even in the presence of short notices many users installed programs that disclosed potentially harmful privacy and security practices, but soon afterwards they regretted their actions. Similarly, Spiekermann et al. (2001) showed that behaviour is not responsive to different disclosures about privacy practices. In a shopping experiment, they presented two different privacy statements to consumers. In one treatment, consumers were reminded of the shelter provided by EU data protection regulations. In the other treatment, consumers were presented with a harsh privacy notice suggesting that data would be sold to a third party and no protection would be provided. The researchers observed no behavioural difference between the treatments. The researchers also found that privacy preferences are a somewhat consistent predictor for revealed privacy behaviours, but that even privacy fundamentalists reveal an unexpectedly large amount of personal information when involved in engaging shopping tasks (Spiekermann et al., 2001).

The academic literature on online notice and consent and the lessons-learned from self-regulatory and governmental efforts to improve the status quo are substantial (Smith et al., 2011), but ultimately inconclusive to our context. Specifically, to the best of our knowledge, no study exists that directly explores the context of post-transaction marketing or negative options practices.

2.3 Opt-In and Opt-Out Options

Our study touches upon several dimensions of computer-human interaction research. In particular, we want to highlight the work on opt-in and opt-out mechanisms. When considering post-transaction marketing offers, an inappropriate opt-in or opt-out design may contribute to consumers making decisions that they later regret.

Different organizations hold different attitudes towards these two approaches. Since the European Union (EU) Data Directive of 1995, European law requires that opt-in choices must be used for consumer protection. The more recent E-Privacy Directive, suggests that information to a user's data can only be gained if the "user concerned has given his or her consent, having been provided with clear and comprehensive information... about the purposes of the processing (European Union, 2009)." However, it is unclear to which degree these regulations apply to post-transaction marketing (in which the consumer has already agreed to share data during a primary transaction). In contrast, industry organizations such as the Direct Marketing Association recommend an opt-out procedure (Direct Marketing Association, 2002).

Previous e-commerce studies have shown opt-in and opt-out mechanisms have a significant impact on users' participation in online activities (Lai and Hui, 2006; Staten and Cate, 2003). From the service providers' point of view, Staten and Cate pointed out that opt-in strategies would raise company's cost and lower profits (Staten and Cate, 2003). Such procedures may also generate more unneeded offers to uninterested or unqualified consumers, and raise the number of missed opportunities for interested consumers. From the consumers' point of view, Lai and Hui studied how frames, defaults, and privacy concerns impact users' online newsletter subscription behaviours. They found that opt-in approaches result in lower levels of participation than opt-out approaches (Lai and Hui, 2006).

To the best of our knowledge, we are unaware of any empirical evaluations of the default opt-in/opt-out settings with respect to post-transaction marketing. While the e-commerce studies on the effects of opt-in and opt-out mechanisms are relevant to our study, we note that in our scenario the resulting

trade-offs are more complex. For legitimate e-commerce activities, the goal is to increase user participation and the amount of personal data being shared. However, from a consumer and regulatory perspective, as well as from the viewpoint of legitimate businesses, the prevalence of misleading marketing practices undermines trust in e-commerce and might have a negative long-term sales impact.

3 Methodology

3.1 Amazon Mechanical Turk

This study was conducted on Amazon Mechanical Turk (AMT), which is a platform that connects requesters of services (e.g., researchers) with individuals willing to perform tasks. Experiments on AMT are able to reach a large number of potential subjects in a relatively short period of time, for a cost comparatively lower than traditional laboratory studies (Nochenson and Grossklags, 2013).

AMT has been used for academic studies ranging from economics (Horton et al., 2011) to cognitive decision-making (Reitter et al., 2013). While AMT was originally intended to perform tasks that were difficult to automate (e.g. translation; see Callison-Burch, 2009), the service has since gained popularity and is commonplace in behavioural research including privacy and security studies. For example, Sheng et al. (2010) investigated susceptibility to email-based phishing schemes, Christin et al. (2012) studied individuals' willingness to engage in unsafe online behaviours in exchange for payments, Wang et al. (2013) and Xu et al. (2012) researched how users engage with privacy configuration interfaces when installing social applications. Further, despite concerns about the validity of using AMT for research studies, it has been shown that AMT participants "produce reliable results consistent with standard decision-making biases" (Goodman et al., 2012).

Our AMT study setup required participants to be from the United States and to have an approval rate of at least 85% for previously completed tasks. These restrictions were put in place in order to ensure quality and relevance to ecommerce practices in Western countries (Mason and Suri, 2012; Ross et al., 2010; Ipeirotis, 2010).

3.2 Experimental Setup and Sequence

For this study, we created an artificial music store entitled MelodiesFor.Us and registered the same domain name. The sequence of the experiment was as follows:

1. On AMT, potential participants could view basic information about the study on "music purchasing behaviors" and were able to "accept" the task. Participants were then presented with a consent form (per the University IRB) detailing, for example, the payment structure. The consent form did not detail the exact procedure that each potential participant was to encounter; since that would have compromised the study (i.e., informing them about the post-transaction offer would render our interventions superfluous).
2. After acceptance, participants were directed to an instruction screen, which explained to participants that they were to enter a music store, and that they would receive \$1.50 to purchase a song (at the cost of \$0.99). They were instructed that "any transaction in the shopping environment reduced your starting budget as described in the shopping environment" and that they would receive any leftover funds from their budgets at conclusion of the experiment.
3. Participants were then directed to the music store (Figure 1). Participants could view and sample six different songs. Participants had to add one song to their cart to continue the experiment.
4. Participants were then shown a checkout page that requested additional information for purchase confirmation. This page requested their AMT ID, age, zip code, and email. Participants then confirmed the purchase of the song.

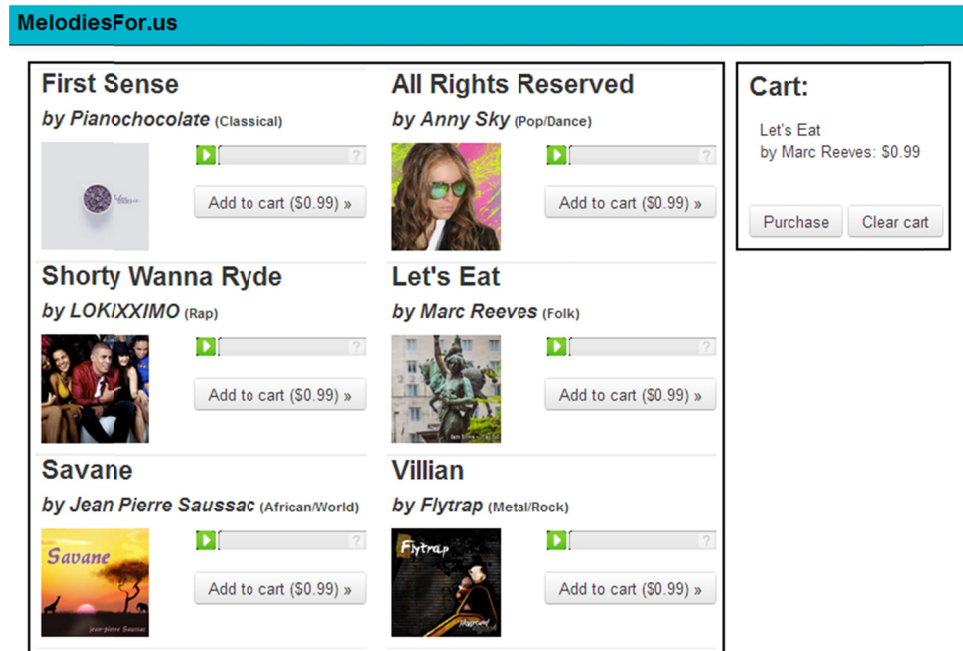


Figure 1. *MelodiesFor.us shopping page. On this page, participants were able to sample the six songs and select a single song for purchase.*



Figure 2. *Post-transaction offer page. This page offered the SafeDelivery product to participants. The box in the bottom left changed based on the treatment condition. The text on the right describes the service. If participants purchased the offer, they were charged \$0.50 of their budget and received an email containing the same song that they had previously downloaded.*

5. Participants were then presented with the post-transaction offer screen offering a product called “SafeDelivery” (Figure 2). If, based on the choices on this screen, a participant purchased the product then their endowment was reduced by \$0.50 leaving \$0.01 of their initial budget.
6. After the post-transaction offer screen, participants were asked to complete a survey.
7. When participants successfully completed the survey, they were shown a task completion message, and were instructed to indicate as such in the AMT interface to receive payment.

3.3 Payment Structure: Participation Fee and Bonus Payment

In accordance with experiments in the experimental economics tradition, participants were compensated with a participation fee and a bonus payment based on their performance (Grossklags, 2007). The participation fee for participants was \$0.50, which they earned for completing all parts of the experiment. Participants that did not complete all parts of the experiment were not compensated. However, participants could end the study at any time, if they did not wish to continue.

For the experiment, participants were given a budget of \$1.50 (separate from the participation fee) to purchase one song for \$0.99. After the music purchase, participants were redirected to the post-transaction offer. If participants did not purchase the offer, they were paid as a bonus their remaining endowment (\$0.51). If they did purchase the post-transaction offer, their budget was reduced by the cost of the service (\$0.50) and they were paid the remaining \$0.01 as a bonus.

3.4 SafeDelivery (Post-Transaction Offer)

SafeDelivery is a mock service that we designed as a post-transaction marketing offer. The layout of the offer page was specifically designed to look similar, but not identical to, the rest of the MelodiesFor.us site. The consistency in interface likely increases the perceived trustworthiness of the post-transaction offer (Edelman, 2009). The offer page includes wording that is somewhat difficult-to-comprehend. It offers a “50% discount” on another copy of the song (i.e., for \$0.50 extra, participants would receive a second copy of initially purchased song). The offer is designed so that the inattentive participant could mistake the discounted additional product for a discount on the original purchase.

3.5 Treatments

We designed six treatment conditions that were all expressed on the post-transaction page and which represent three different design choices which could be mandated by law or self-regulation (and are partly responsive to ROSCA). The treatments are summarized in Table 1.

Treatment	Opt-in/out	Email pre-populated?
T1	In	Yes
T2	Out	Yes
T3	In	No
T4	Out	No
T5	In	Yes-hidden
T6	Out	Yes-hidden

Table 1. *Treatments and their properties.*

As a first design choice, we test the impact of post-transaction offers that are opt-in against those that are opt-out. That is, the treatments were partitioned into half opt-in and half opt-out (i.e., three treatments each). In treatments that were opt-in, participants were required to hit a large button to opt into the purchase of SafeDelivery (costing \$0.50), while in those that were opt-out, participants were required to hit the same large button (with different text) to opt out of the purchase; see Figure 3.

As a second design choice, we are interested in how the use of data pass techniques affects the number of participants who purchase the post-transaction offer. To investigate this, treatments varied by the information that needed to be entered to purchase the post-transaction offer. In a subset of treatments, users were required to enter an email address (with valid syntax) in order to hit the button, while in other treatments the email address box was pre-populated with the participant's email address. The data was derived from previously entered information during the checkout phase in the music store which simulated a data pass arrangement between the primary merchant and the post-transaction marketer.

The third design choice that we are investigating is the impact of visualizing to the participant that data pass is taking place or not (i.e., its visibility). That is, we are interested whether there is a noticeable impact on consumer behaviour if they are aware of their data being passed from the primary merchant to the post-transaction marketer. To test this effect, in some treatments participants were not shown the email address box (while the information was transferred in the background).

Depending on the treatment, the disclosure text in the large box (see Figure 2) was slightly customized.

Entering your email address below constitutes your electronic signature and we will automatically charge your budget according to the Offer Details to the right.

E-mail Address:

Email song

[No thanks](#)

You are subscribed to the SafeDelivery service. Entering your email address below constitutes your electronic signature and we will remove you from the SafeDelivery service described to the right.

E-mail Address:

Remove me from SafeDelivery

[No thanks](#)

Figure 3. Visual display of opt-out and opt-in treatments (left is T3 and right is T4).

3.6 Post-Experimental Survey

After visiting the post-transaction offer page and either purchasing the offer or not, participants were directed to the post-experimental survey including questions about basic demographic information (age, education, gender), questions about the post-transaction offer they had just seen, questions encouraging individual reflections on the experience of participating in this experiment, questions about shopping behaviours, and attitudes towards privacy and risk avoidance online, as well as questions from the Social Comfort section of the Internet Use scale in Davis et al. (2002). Additionally, the survey asked a number of integrity questions that served to ensure participants filled out the survey in a mindful way. The types of questions that we call integrity questions are also sometimes referred to as “screeners” in the language of Berinsky et al. (2012). These questions aim to ensure that participants are paying attention during the survey.

4 Results

4.1 Attrition and Privacy Choice

In total, 742 individuals began the study. Of this group, 550 (74%) participants completed the study. As in a typical electronic commerce situation, we expected that many participants would not complete our purchasing study due to attrition (Sismeiro and Bucklin, 2004).

Particularly relevant for our study, on the “checkout” page participants had to enter their email address, zip code, age, and AMT ID. 10% of all participants apparently chose not to provide this information, and instead left the task (Table 2). Presumably, the solicitation of their information caused participants to make a conscious privacy choice.

Last page	# of Participants (%)
Consent form	4 (0.5%)
Instructions	2 (0.3%)
Music shopping	15 (2%)
Checkout song	77 (10%)
Post transaction offer	56 (8%)
Post survey	26 (4%)
Finished	550 (74%)

Table 2. Attrition rates by page.

Treatment	# Drop on Post-transaction offer page (% of total drops)
T1	8 (4.1%)
T2	9 (4.7%)
T3	5 (2.6%)
T4	17 (8.9%)
T5	7 (3.7%)
T6	10 (5.2%)

Table 3. Attrition rates by treatment on post-transaction offer page.

We identified 12 participants (1.62%) that continued to the last page of the experiment but did not complete it correctly due to any number of potential technical difficulties. Those subjects have been removed from the analysis.

4.1.1 Attrition by Treatment

We examined attrition differences across treatments. We focus on the post-transaction offer page, since its design varied by treatment (Table 3). Participants in T3 (opt-in, not-pre-pop) had the lowest attrition rate, while participants in T4 (opt-out, no-pre-pop) had the highest. Since T3 is the least aggressive treatment (because it requires the most conscious intervention and is opt-in), it is unsurprising that it has the lowest attrition. Conversely, T4 is the most aggressive treatment (since it is opt-out and requires the most effortful action to opt-out) and has the most attrition.

Attrition rates were also lower for treatments that were opt-in (T1, T3, T5) than for their opt-out counterparts (T2, T4, and T6, respectively). Since opt-in is the preferred way for websites to behave (as defined by policy makers and consumer advocates, and the legal way according to ROSCA), this result is not surprising. When faced with a choice that is unusual or potentially untrustworthy, it makes sense that participants are more likely to abandon the purchasing episode. Paradoxically, in practice, participants would remain signed up for the service since they did not actively opt-out.

In the following subsections, analysis of this experiment will be restricted to those participants that successfully completed the experiment (i.e., 550 participants).

4.2 Demographic Information

For the 550 participants that finished the experiment, the mean age was 33 years old ($\sigma = 11$ years). 292 participants (53%) were female. Most participants had gained at least some college education. This demographic mix of participants was consistent with previously-completed surveys on the demographics of participants on AMT (Ipeirotis, 2010; Horton et al., 2011).

4.3 Value of Post-Transaction Offer

To address the problem area of potentially unwanted and deceptive post-transaction marketing, we aimed to simulate a situation where the post-transaction offer is of little or no value to participants. Therefore, the act of reducing the number of conversions can be understood as a socially optimal act.

To evaluate our study design, we investigated the perceptions of participants regarding the offer. As expected, 235 participants (43%) rated the offer's value the lowest possible value on a 5-point Likert scale, and only 35 participants (6%) rated it the highest. The median value was 2.

We found a strong linear relationship between how participants perceived the products' value and whether they purchased it. Participants that rated the product very valuable (rating of 5) purchased it 77% of the time, whereas participants that rated the product as a 3 purchased it 55% of the time, and those that rated the product a 1 purchased it 24% of the time. It is surprising that over 20% of the participants that thought the service was not at all valuable still purchased it. Similarly, Good et al. (2005 and 2007) found that installation decisions of potentially harmful programs were regretted immediately after installation and further reflection.

4.4 Conversion Rates across all Treatments

Treatments took different approaches to reduce the number of participants who purchased the post-transaction offer (assuming that a purchase would be an undesirable outcome). Of the 550 participants that finished the experiment successfully, 222 participants (40%) purchased the post-transaction offer (across treatments). Treatments do have significantly different conversion rates ($\chi^2 = 51.6, p < 0.0001$); see Figure 4 and Table 4.

Treatment	Opt-in/out	Email pre-populated?	% Conversions (i.e., Email sent) (# convert/# total)
T1	In	Yes	45.3% (43/95)
T2	Out	Yes	44.4% (40/90)
T3	In	No	18.7% (17/91)
T4	Out	No	68.6% (59/86)
T5	In	Yes-hidden	32.6% (30/92)
T6	Out	Yes-hidden	34.4% (33/96)

Table 4. Treatment details and their conversion rates.

T3's conversion rate is significantly lower than all other treatments ($\chi^2 = 20.2, p < 0.0001$). This result was expected, since it featured the most effortful procedure to purchase the offered product. In order to agree to the post-transaction offer in this treatment, participants were required to manually enter their email address and hit the large button to confirm the purchase.

The result is contrasted by a treatment with the perhaps most unexpected and potentially misleading procedure, i.e., T4. The treatment's conversion rate is significantly higher than all other treatments ($\chi^2 = 32.4, p < 0.0001$), because it required an increased effort to not purchase the offer. In this treatment, participants were set to agree to the post-transaction offer by default. To change this outcome, they were required to enter their email address and click the large button. Nearly 70% of the participants did not take these two actions, which left them with the default outcome. (Note that a popup box would notify the participants if they merely clicked the big button without entering their email address.) The size of the effect is perhaps surprising. However, we expected the general tendency of the result due to the increased effort to opt-out, and the fact that clicking "No thanks" means to stay with the default outcome (as stated in the disclosures on the post-transaction offer screen). The latter aspect was likely unexpected for participants who might have mistakenly presumed that "No thanks" would result in a cancellation of the purchase. This behaviour is, therefore, likely attributable to optimism biases and conditioned-response biases (Meyer, 2009).

4.5 Opt-In versus Opt-Out

When determining the effectiveness of opt-in and opt-out treatments, it is only reasonable to compare pairs of treatments that have all the same properties except for opt-in or opt-out status. Therefore, the pre-populated treatments T1 and T2 can be compared, the pre-populated and hidden information

treatments T5 and T6 can be contrasted, and the not pre-populated treatments T3 and T4 can be compared.




Figure 4. Conversion rates within treatments. T3 has the least conversions and T4 has the most. There is no difference between T1 and T2, or between T5 and T6. T1+2 have more conversions than T5+6. Bars represent 95% confidence intervals.

In the case where the user's email is pre-populated on the offer page (i.e., T1 and T2), there is no significant difference in conversion between participants who were shown offers with the opt-in option and those with the opt-out option. In the case where the user's email is pre-populated and hidden (i.e., T5 and T6), there is also no significant difference in conversion between participants that were presented with opt-in and opt-out offers. When participants' email addresses were not pre-filled, the opt-out treatment T4 had a significantly higher conversion rate than its opt-in counterpart T3 ($T4 > T3$; $\chi^2 = 45, p < 0.0001$). See our discussion in Section 4.4.

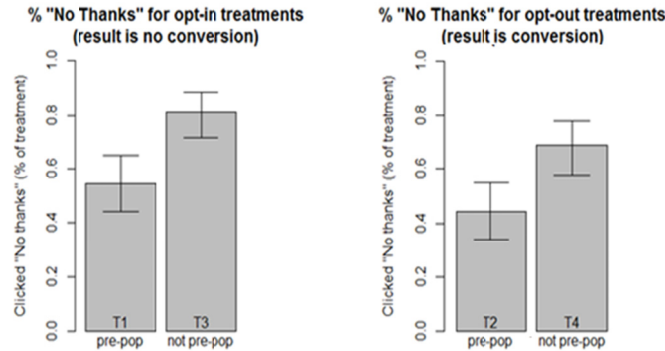


Figure 5. Number of occurrences of clicking "No thanks" by treatment and split by opt-in/opt-out. In opt-in treatments (left), clicking "No thanks" resulted in a non-conversion (declined offer), while in opt-out treatments (right), clicking "No thanks" resulted in a conversion. Bars represent 95% confidence intervals.

4.6 Effectiveness of Data Pass

We can investigate the impact of data pass arrangements from two perspectives. First, we can compare the rate of conversion between groups of participants that were presented with offers that utilized data pass and those that did not. Second, we can also investigate how often participants selected the "No

thanks” button, presumably to cancel the post-transaction marketing purchase (even if it actually resulted in the opposite outcome). Note that these metrics are not complete mirror images. For opt-in treatments T1 and T3, “No thanks” is indeed the opposite of the conversion rate. However, for opt-out treatments T2 and T4, conversion rate and the number of times participants selected “No thanks” are the same. The two different perspectives make it easier to visually uncover the impact of conditioned response biases (Meyer, 2009), even if the underlying conversion rates are of course unchanged.

Using the metric of conversion rate, experimental results show that data pass (simulated by a pre-filled box on the post-transaction offer page) increases the conversion rate for treatments that present the offer as opt-in and decreases the conversion rate for treatments that present the offer as opt-out. Treatment T1, which presents an opt-in offer using data pass, has a conversion rate which is more than twice as high as that of Treatment T3, which presents an opt-in offer without data pass (this relationship is significant; $T1 > T3$; $\chi^2 = 15.0, p < 0.001$). Treatment T4, which presents the post-transaction offer using data pass and an opt-out strategy, has a conversion rate over 50% higher than Treatment T2, which uses opt-out and does not use data pass (this relationship is significant; $T4 > T2$; $\chi^2 = 10.4, p < 0.01$).

Using the metric of “No thanks” rate, experimental results show that data pass leads fewer participants to select default options (Figure 5). In Treatment T1, which presents an opt-in offer using data pass, fewer people select “No thanks” compared to Treatment T3, which presents an opt-in offer without data pass (this relationship is significant; $T1 < T3$; $\chi^2 = 13.8, p < 0.001$). Also in treatment T4, which presents the post-transaction offer using data pass and an opt-out strategy, fewer people select “No thanks” compared to Treatment T2, which uses opt-out and does not use data pass (this relationship is significant; $T2 < T4$; $\chi^2 = 9.5, p < 0.01$).

The extra effort required to type an email address into a box reduces the number of participants who take this more effortful option regardless if that option is to opt in or to opt out. Likewise, by making it apparently less effortful to choose the “No thanks” option, consumers will do so, even to their own detriment.

4.7 Consumer Knowledge of Data-Pass Arrangements

The pre-populated and hidden information treatments T5 and T6 offer a different perspective on the privacy impact of post-transaction marketing. These treatments appear less like a privacy choice, since there is no active solicitation of information (i.e., of the participant’s email address) nor is there a clear indication that information is being transferred (i.e., a pre-populated box with personal data). Treatments T5 and T6, which do not disclose the use of data pass, have conversion rates that are around 25% lower than their counterpart treatments T1 and T2 (with the same opt-in and opt-out status), respectively. However, this difference is not significant; presumably due to lack of power.

5 Discussion and Concluding Remarks

In this experiment, we tested consumers’ susceptibility to fall for costly post-transaction marketing offers that in return yield little-to-no value to the consumer. Of the 550 participants, over 40% signed up for the SafeDelivery product giving them access to a song (for an additional payment), that they had just purchased and downloaded a few seconds earlier in the MelodiesFor.Us store. It is difficult to argue that these purchases were wanted, and in fact in the post-experimental survey most consumers agreed that the service is not valuable. The monetary consequences of these unwanted purchases were considerable (i.e., 10%-40% of the average hourly wage of a typical AMT worker; Horton and Chilton, 2010).

In our online experiment, we researched several (but not all) design characteristics of post-transaction marketing (Cox, 2010; U.S. Senate, 2009). We used a design for the post-transaction marketing offer that is derived from legal cases (e.g. the Intelius case; see Johnson 2010). Then, we systematically

varied the data pass arrangements and its visibility, and the process of gaining user consent (e.g., opt-in/opt-out). We kept other user interface aspects constant across our experimental treatments.

We collected sizeable samples for six experimental treatments that help to inform public policy and self-regulatory efforts by marketers and online businesses. For example, in the United States, the Restore Online Shoppers' Confidence Act (ROSCA) likely outlaws the opt-out treatments that were part of this experiment (i.e., Treatments 2, 4, and 6). Given our results, this policy direction appears partially effective. In each opt-out treatment, the same number of consumers or more purchased the post-transaction service compared to its opt-in counterpart.

ROSCA also requires that the terms of the post-transaction offer be present before purchase. In our view, this requirement is insufficient. In this experiment, we did not omit material facts from the participants. However, the post-transaction offer was written in positive language, and pragmatically implied that SafeDelivery is a particularly useful service. In many jurisdictions, such "puffery" is usually discounted since consumers are thought to be able to identify exaggerations. We are unconvinced of this claim, in particular, when consumers have already completed a primary transaction and as a result are more likely to have positive perceptions about the shopping experience.

Despite the visual dissimilarities between the shop/check-out pages and the post-transaction offer (and the order confirmation shown at the top of the page), many consumers had difficulty realizing that a second distinct transaction is taking place. For example, from our post-experimental survey, we learned that about 50% of the participants believed that SafeDelivery is the same company as MelodiesFor.Us, or they were unable to tell afterwards.

ROSCA's mandate that requires an effortful confirmation is moderately effective at reducing unwanted purchasing of post-transaction offers. The effect of this policy change is more substantial than the effect we observed from spending more time on the post-transaction marketing page (which presumably increases the likelihood of scrutinizing the offer in more detail).

However, all effortful confirmations are not created equally. On the one hand, we found that participants that engaged with the most effortful confirmation to purchase a post-transaction offer (i.e., typing in their email and clicking a button) were the least likely to purchase the product. On the other hand, participants that needed to type in an email address to opt-out from the purchase were the least likely to do so, (again) due to the increased effort required. It is, therefore, imperative to pay special attention to opt-out regimes that require consumer effort.

Surprisingly, participants that were presented with a pre-populated offer page were more likely to sign up for the service compared to the treatment where the email information was hidden from view. We believe that the visual presence of the email address served to reinforce the belief that the primary transaction was still in process. Such misleading inferences are likely a result of schematic thinking. I.e., the unusual interaction with the post-transaction marketing offer was insufficient to disrupt the schematic belief of being part of a "standard ecommerce" transaction (Alba & Hasher 1983; Harris & Monaco 1978). We also offer a different view. Consumers may interpret the post-transaction offer as a combination of a privacy sale (i.e., giving away their email address) and the SafeDelivery product. It is conceivable that the privacy aspect of the transaction distracts participants from a careful evaluation of the post-transaction offer page. Given the lower sign-up rate of the treatment where the email address is hidden, it might be prudent to recommend that privacy transactions should be treated separately from other transactional aspects.

Acknowledgement

We thank the anonymous reviewers and the associate editors for their comments on an earlier version of this paper. We further gratefully acknowledge the support given by a Google Faculty Research Award and a College of Information Sciences and Technology seed grant at the Pennsylvania State University.

References

- Acquisti, A., and Grossklags, J. (2005). Privacy and Rationality in Individual Decision Making, *IEEE Security and Privacy*, 3 (1), 26-33.
- Alba, J., and Hasher, L. (1983). Is memory schematic? *Psychological Bulletin*, 93 (2), 203-231.
- Ballaré, D., and Von Bergen, C. (2008). Negative option marketing and ethical theory. *International Academy of Business and Public Administration Disciplines*, Dallas, Texas, April 24-27.
- Boush, D., Friestad, M., and P. Wright, P. (2009). *Deception in the marketplace: The psychology of deceptive persuasion and consumer self protection*. Routledge, New York, NY.
- Callison-Burch, C. (2009). Fast, cheap, and creative: Evaluating translation quality using Amazon's Mechanical Turk. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP '09)*, Association for Computational Linguistics, pp. 286-295.
- Christin, N., Egelman, S., Vidas, T., and Grossklags, J. (2011). It's all about the Benjamins: An empirical study on incentivizing users to ignore security advice. In *Proceedings of the Fifteenth International Conference on Financial Cryptography and Data Security*, Gros Islet, St. Lucia, pp. 16-30.
- Cox, P. (2010). The Invisible Hand of Preacquired Account Marketing. *Harvard Journal on Legislation*, Vol. 47, No. 2.
- Direct Marketing Association (2002). *Direct Marketing Association's Online Marketing Guidelines and Do the Right Thing Commentary*. Available at: <http://www.the-dma.org/guidelines/onlineguidelines.shtml>.
- Edelman, B. (2009). Prepared Statement of Benjamin Edelman presented to the United States Senate Committee on Commerce, Science, and Transportation. <http://www.benedelman.org/posttransaction/statement-1.pdf>.
- European Union (1995). Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of such Data. Available at: <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:en:HTML>.
- European Union (2009). Directive 2009/136/EC of the European Parliament and of the Council (E-Privacy Directive). Available at: <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:337:0011:0036:En:PDF>.
- Federal Trade Commission (1984). *FTC Policy Statement on Deception*. <http://www.ftc.gov/ftc-policy-statement-on-deception>.
- Good, N., Dhamija, R., Grossklags, J., Aronovitz, S., Thaw, D., Mulligan, D., and Konstan, J. (2005). Stopping Spyware at the Gate: A User Study of Privacy, Notice and Spyware. In *Proceedings of the Symposium On Usable Privacy and Security*, Pittsburgh, PA, pp. 43-52.
- Good, N., Grossklags, J., Mulligan, D., and Konstan, J. (2007). Noticing Notice: A large-scale experiment on the timing of software license agreement. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, San Jose, CA, pp. 607-616.
- Goodman, J., Cryder, C., and Cheema, A. (2013). Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making*, 26 (3), 213-224.
- Grossklags, J. (2007). *Experimental Economics and Experimental Computer Science: A Survey*, Workshop on Experimental Computer Science (ExpCS'07), San Diego, CA.
- Harris, R., and Monaco, G. (1978). Psychology of Pragmatic Implication: Information Processing Between the Lines. *Journal of Experimental Psychology: General*, 107 (1), 1-22.
- Horton, J., and Chilton. L. (2010). The labor economics of paid crowdsourcing. In *Proceedings of the 11th ACM Conference on Electronic commerce*, New York, NY, pp. 209-218.
- Horton, J., Rand, D., and Zeckhauser, R. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics* 14 (3), 399-425.
- Ipeirotis, P. (2010). *Demographics of Mechanical Turk*. NYU Working Paper No. CEDER-10-01. Available at SSRN: <http://ssrn.com/abstract=1585030>.

- Jensen, C., and Potts, C. (2004). Privacy policies as decision-making tools: An evaluation of online privacy notices. In *Proceedings of SIGCHI Conference on Human Factors in Computing Systems*, Vienna, Austria, pp. 471-478.
- Johnson, G. (2010). Intelius to pay \$1.3M for deceptive marketing. *The Seattle Times*. http://seattletimes.com/html/localnews/2012587904_apwainteliussettlement.html.
- Lai, Y., and Hui, K. (2006). Internet opt-in and opt-out: Investigating the roles of frames, defaults and privacy concerns. In *Proceedings of the 2006 ACM SIGMIS CPR Conference on Computer Personnel Research*, Pomona, CA, pp. 253-263.
- Mason, W. and Suri, S. (2012). Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods*, 44 (1), 1-23.
- Meyer, R. (2009). Prepared Statement of Robert J. Meyer presented to the United States Senate Committee on Commerce and Transportation.
- Milkman, K., Chugh, D., and Bazerman, M. (2009). How can decision making be improved? *Perspectives on Psychological Science*, 4 (4), 379-383.
- Nochenson, A., and Grossklags, J. (2013). A Behavioral Investigation of the FlipIt Game. In *Proceedings of the Twelfth Workshop on the Economics of Information Security (WEIS 2013)*, Washington, DC.
- Reitter, D., Grossklags, J., and Nochenson, A. (2013). Risk-Seeking in a Continuous Game of Timing. In *Proceedings of the Twelfth International Conference on Cognitive Modelling (ICCM)*, Ottawa, Canada, pp. 397-403.
- Ross, J., Irani, L., Silberman, M., Zaldivar, A., and Tomlinson, B. (2010). Who are the crowdworkers?: Shifting demographics in Mechanical Turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Atlanta, GA, pp. 2863-2872.
- Sheng, S., Holbrook, M., Kumaraguru, P., Cranor, L., and Downs, J. (2010). Who falls for phish?: A demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Atlanta, GA, pp. 373-382.
- Sismeiro, C., and Bucklin, R. (2004) Modeling Purchase Behavior at an E-Commerce Web Site: A Task-Completion Approach. *Journal of Marketing Research*, 41 (3), 306-323.
- Spiekermann, S., Grossklags, J., and Berendt, B. (2001). E-privacy in 2nd generation E-Commerce: Privacy preferences versus actual behavior. In *Proceedings of the Third ACM Conference on Electronic Commerce*, Tampa, Florida, pp. 38-47.
- Smith, H., Dinev, T., and Xu, H. (2011) Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35 (4), 989-1016.
- Staten, M., and Cate, F. (2003). The Impact of Opt-In Privacy Rules on Retail Credit Markets: A Case Study of MBNA. *Duke Law Journal*. 52 (4), 745-786.
- Tufte, E. (1997). *Visual explanation: Images and quantities, evidence and narrative*, Graphics Press, Cheshire, CT.
- Turow, J., Hoofnagle, C., Mulligan, D., Good, N., and Grossklags, J. (2007). The Federal Trade Commission and Consumer Privacy in the Coming Decade. *I/S: A Journal of Law and Policy for the Information Society*, 3 (3), 723-749.
- U.S. House (2010). 111th Congress, 2nd Session. S 3386 RFH, Restore Online Shoppers' Confidence Act. Washington, Government Printing Office, 2010.
- U.S. Senate (2009). Staff Report for Chairman Rockefeller, Aggressive Sales Tactics on the Internet and their impact on American Consumers, Committee on Commerce, Science, and Transportation, November 16, 2009. Available at: <http://www.gpo.gov/fdsys/pkg/CHRG-111shrg54917/pdf/CHRG-111shrg54917.pdf>.
- Wang, N., Grossklags, J., and Xu, H. (2013). An Online Experiment of Privacy Authorization Dialogues for Social Applications. In *Proceedings of the 16th ACM Conference on Computer Supported Cooperative Work and Social Computing*. San Antonio, Texas, pp. 261-272.
- Xu, H., Wang, N., and Grossklags, J. (2012). Privacy by ReDesign: Alleviating Privacy Concerns for Third-Party Apps. In *Proceedings of the International Conference on Information Systems*, Orlando, FL.