
Too Sensitive to Share? Working with Consumers' Credit Card Transactions

Johannes Huebner

ETH Zurich
Zurich, Switzerland
jhuebner@ethz.ch

Alexander Illic

ETH Zurich
Zurich, Switzerland
ailic@ethz.ch

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI'18 Extended Abstracts, April 21–26, 2018, Montreal, QC, Canada

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5621-3/18/04...\$15.00

<https://doi.org/10.1145/3170427.3174350>

Abstract

Simply by paying with credit cards, consumers spend more than they would with cash due a difference in payment transparency. We introduced a mobile application to test the efficacy of personalized feedback interventions to help people save money by lowering credit card expenses, and ultimately, to guide them towards a more responsibly use of digital forms of payment. For our large-scale field study (N>1'000 individuals), we cooperated with a credit card issuer to be able to test the effectiveness of our mobile-mediated interventions on real-world credit card transactions over a period of three months. This paper summarizes the main challenges we encountered, and the taken measures that enabled us to leverage highly sensitive data for research, which serve as guidelines for future industry-facing field studies.

Author Keywords

Industry Research; Sensitive Data; Credit Card Transaction Salience; Financial Literacy; Field Study

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous

Introduction

In many economies, we observe a successive replacement of cash by digital payment forms, such as debit or credit cards [1]. Such cashless means of payment offer a number of advantages to consumers, such as added convenience, greater access to money (i.e. liquidity, in the case of credit cards), as well as the possibility to track expenses post-hoc, allowing consumers to learn from and adjust their spending behavior. However, research suggests that these benefits come at a price. Consumers spend more money on the same items, focus more on product properties rather than the associated cost, and are more likely to indulge in treats and luxury items when using credit cards instead of cash [2,6,8,9,11]. Thus, the choice of payment form imposes a trade-off on consumers, between convenience and post-hoc transparency when using cashless forms of payment, and financial discipline when paying in cash.

In our study, we sought to address this trade-off by using mobile information systems to support people when making cashless payments. For this purpose, we developed a mobile application for iOS and Android (see Figure 1) phones to test the effect of increased transaction salience on consumers' credit card spending behavior.

Thanks to a collaboration with a large Swiss credit card issuer, we had the unique opportunity to conduct research with a sizeable group of consumers ($N > 1'000$) using their credit cards in the real world whilst receiving feedback through our app. In this paper, we share our key learnings with the academic community to facilitate future practice-oriented research projects.

Related Work

Pain of Payment

Behavioral economists argue that the phenomenon of consumers exhibiting different behavior with credit cards than with cash can be explained with varying intensities of a psychological pain of payment, which depends on the transparency of the used payment form [7,11]. For example, both the form and amount of a payment is much more transparent for cash than for credit cards or special-purpose prepaid cards (e.g. prepaid public transportation cards) [9]. Consumers usually experience two emotions when making a purchase: Pleasure derived from consumption, and pain of paying [7]. When using cash, those emotions are tightly coupled, whereas the pain of payment is extenuated when payments are made digitally [5,8–11]. Credit cards take on a special role, because the actual wealth depletion event occurs with a considerable delay, and because purchases are aggregated into a monthly bill. This makes it challenging for consumers to intuitively learn from purchases and adapt their behavior.

The Role of Exceptional Purchases

In addition to behavioral changes arising from the mere use of digital channels, behavioral economists have long argued that several mental biases can lead to sub-optimal decision-making with costly consequences [14]. Research conducted by [13] highlights the importance of exceptional purchases in short-term financial decision making. They argue that consumers track expenses either explicitly, or implicitly using mental budgets. Research further shows that prior purchases in a particular category deplete the budget of said category, thus rendering the consumer more price-sensitive in future shopping events [12]. Even for

frequently occurring, ordinary expenses like groceries or transportation, consumers often face difficulties in accurately predicting and managing expenditures. However, for less frequently occurring, exceptional purchases like gifts or electronic gadgets, planning errors are more pronounced, arguably because people regard such exceptional purchases in a very isolated fashion (narrow choice bracketing), underestimate their frequency, cut themselves greater financial slack, and thus ultimately overspend [13].

As illustrated above, researchers have identified psychological fallacies as well as behavioral changes arising from the use of information systems as potential sources of sub-optimal decision making in the financial context, for which remedies should be developed.

Method and Challenges

Building on this research, we thus designed a field study following the randomized controlled trial approach. We asked participants to define a weekly spending goal and subsequently manipulated transaction salience by introducing different variants of weekly goal attainment feedback interventions, focused on either ordinary, exceptional, or both types of purchases, whereas a control group received more general goal attainment feedback. We measured the effect on two dependent variables: i) actual credit card spending (relative to the individual baseline) and ii) perceived control over credit card expenditure.

Mobile Application and Experiment Design

Our industry partner, a Swiss credit card issuer, sent out email invitations to an eligible subset of their existing end customers prompting people to download the app and participate in our study over a time frame

of three months. The email invite also included information about the research project.

Inside the app, people were first asked to enter an individual user token from the invitation email, and they were asked for their consent to participate in the study. After opting in, they were randomly allocated to either the Control group or one of three treatment groups, and received group-specific app usage guidelines. Next, an initial questionnaire comprising 19 items surveyed participants for several demographic data points (such as household size and income class), economic attitudes (frugality, compulsiveness), an estimated share of wallet (i.e., which share of all purchases are made with the linked card), as well as financial behaviors (e.g. existence of a household budget). All survey items were implemented as buttons with pre-defined answers, not requiring text input. In a pre-study cohort with N=95 individuals [4], we found that most users filled out the survey in between 70 and 229 seconds (10th and 90th percentile). While a user was taking the survey, our partner firm provided us with a year worth of baseline data (cumulative weekly credit card spending for the prior 52 weeks of the respective customer) through an application programming interface (API, see Figure 2). This enabled us to display the user's mean weekly spending in the following step, which concluded the initial setup process: Participants had to define a weekly spending goal, which could be adjusted during the study. A default goal of spending 15% less than the baseline expenditure was suggested.

Throughout the study, transactions made with the linked credit card were forwarded through the same API, and then displayed to the user in a column chart. Participants in the treatment groups had to categorize

each transaction in a binary schema as either ordinary or exceptional. Finally, each participant received goal attainment feedback as push notification as well as in the app, the content of which depended on the respective condition to which they were randomly assigned in the beginning of the study. After 14 weeks, users were prompted to fill out a final survey, eliciting feedback regarding the mobile application, perceived control over their credit card, as well as changes in their credit card usage habits.

Experimental Method

Even though many studies investigating consumers' financial and shopping behavior use more controlled settings such as lab environments, we opted for a quasi-experimental setup in the real world, because we expected more insightful data by observing consumers' natural behavior over a sufficient amount of feedback cycles. We did not make this choice lightly, however. In addition to the common issues that field studies exhibit, the domain of our research added to the challenges:

- First, while developing a mobile application for research purposes is relatively effortless and affordable, the issuance and distribution of new credit cards solely for research purposes is prohibitively expensive and would require an IT infrastructure and operational setup arguably exceeding the capacity of most research institutions.
- Second, financial transaction data is highly sensitive to both consumers and the involved financial institutions. Even in cases when financial service providers that are not legally prohibited from sharing

their customers' data, they will be very careful in forwarding any data outside organizational boundaries for various reasons, such as potential backlashes from consumers or other stakeholders¹.

- Furthermore, we were subject to the incomplete information problem, since study participants could change their habits and increasingly use cash or other unobserved payment means.
- Finally, both the within-subject as well as the between-subject variance in credit card spending data is rather high, especially when regarded in short intervals, such as on a weekly basis.

In spite of these hurdles, we believe that a field study was the right choice in our case, and we would like to share our key learnings, hoping to thereby contribute to facilitating future industry-oriented research.

Findings and Discussion

Industry Research with Highly Sensitive Data

The previous chapter highlighted some of the challenges we encountered. First of all, the provision of own credit cards appeared disproportionately expensive, hence the only feasible approach was to cooperate with an existing credit card issuer. However, as is also reflected in the rather small number of consumer-facing field studies in the financial domain, financial service providers naturally tend to be cautious about sharing data with third parties due to the sensitive nature of the involved data.

might offer researchers an interesting new way to access consumers' bank transaction histories and account balances.

¹ The EU Directive 2015/2366, more commonly known as Payment Services Directive 2, requires EU-based banks to offer several APIs to their customers' data. Once fully implemented, this

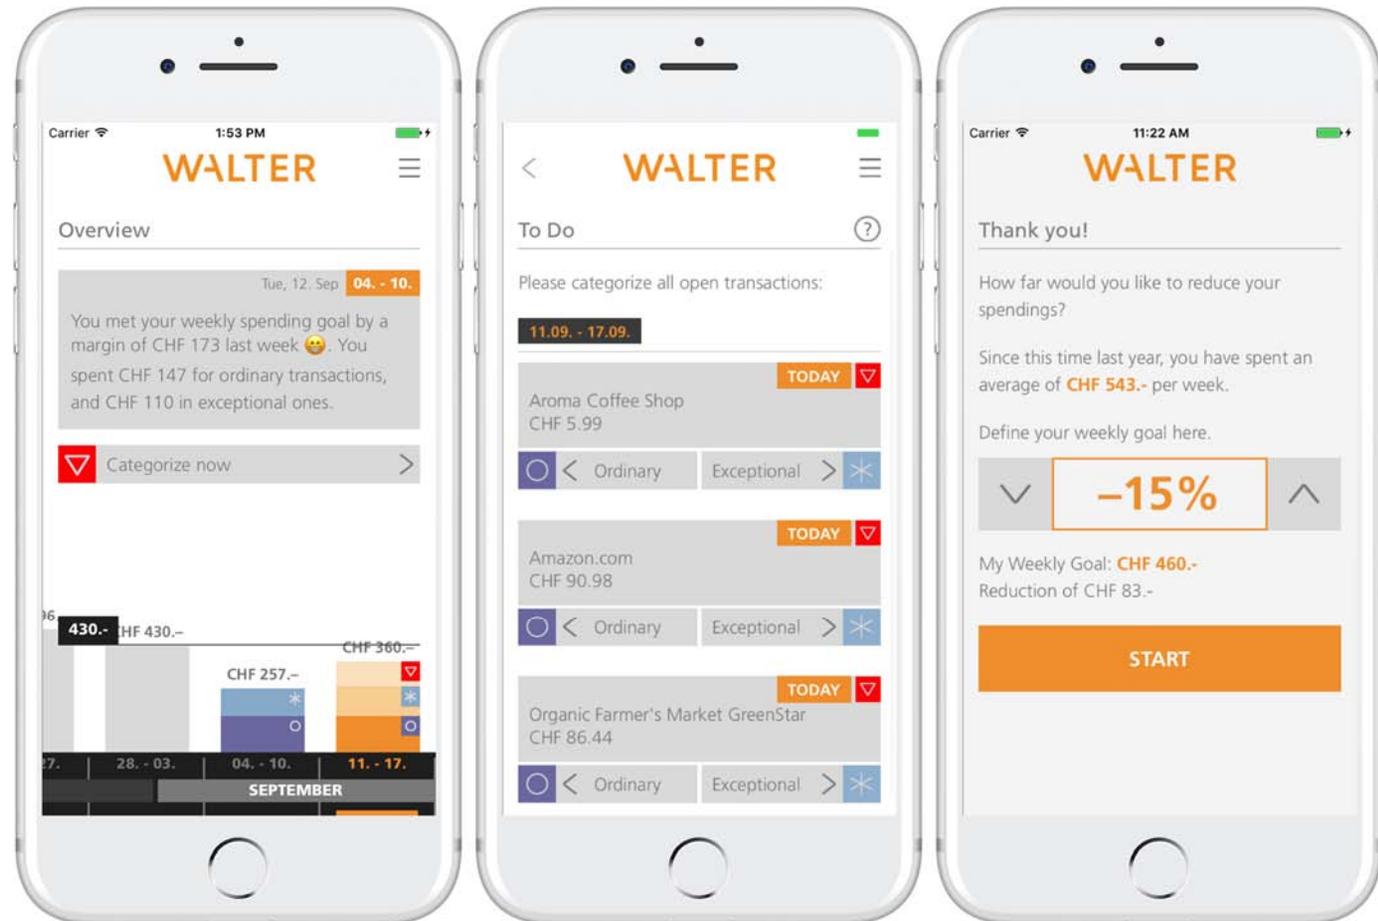


Figure 1: Mobile App Screenshots. Dashboard with column chart displaying the user's weekly expenses, whereas the different colors in a column indicate the categories ordinary, exceptional, or uncategorized expenses (left); Categorization screen where transactions are classified as either ordinary as exceptional (middle); Goal-setting screen, which displays the user's mean weekly cumulative spending over the previous year and suggests a default goal implying a 15% reduction in spending compared to the baseline (right).

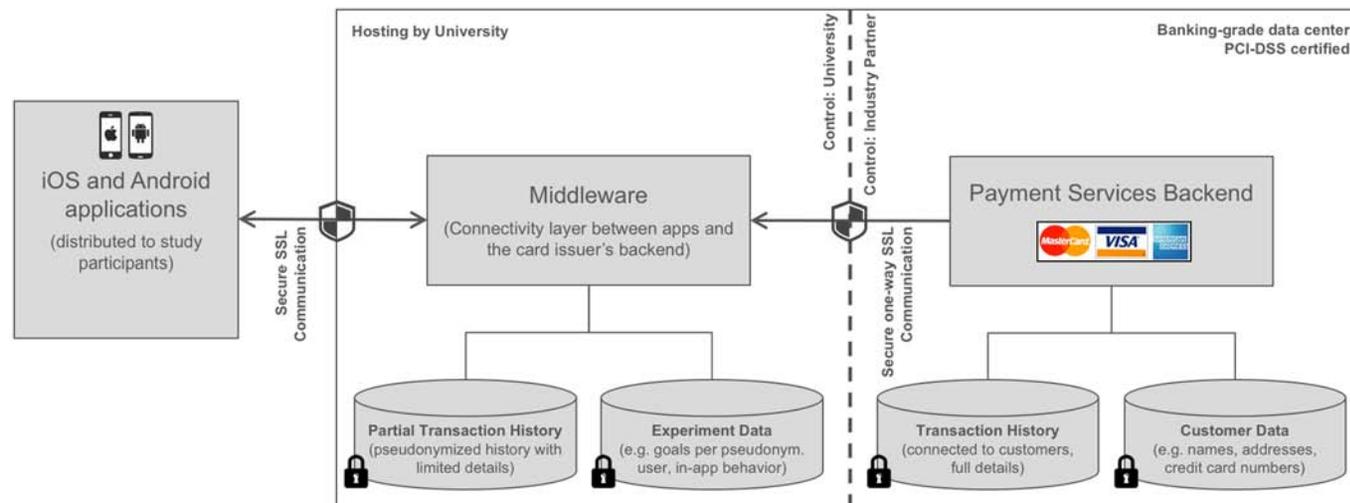


Figure 2: IT architecture implemented for our project with unidirectional data flows from the card issuer's backend to our middleware, and strict separation of data ownership: Any person-identifying data was only known to the credit card firm; transactions were shared with limited details, and in pseudonymized form. Sensitive experiment data, such as economic attitudes or financial behaviors reported by participants in surveys or gathered through in-app analytics, were only accessible to the research team and not shared with the card issuer for privacy reasons.

It was thus essential for us to first establish a trustful relationship as well as a sound legal foundation and data flow setup, which protected the interests of consumers, the credit card firm, and us. We extensively involved the legal teams of both parties, as well as our university's ethics board, especially concerning data ownership and data flows, discussing on a field level which data would be accessible by whom. In addition, the following steps were key to our project's success:

- Any financial data was provided to us in a pseudonymized form, i.e. person-identifying data such as names, email addresses, or credit card numbers, were never shared with our research team.

Instead, the card issuer generated user-specific tokens and linked them with the credit cards in their IT systems. The card firm then sent out email invites to the selected customers, including said tokens (both the selection criteria and email contents had been defined jointly). Each user then had to enter their personal token in the app, which enabled us to deliver their card transactions to them through the app, even though the mapping between tokens and person-identifying information was unknown to us.

- All experiment data including survey data was under the research team's control and was never shared on an individual level to ensure the study participants' privacy with regard to the credit card firm.

- The communication between the credit card firm's IT systems and our own middleware was strictly unidirectional (see Figure 2), i.e. a firewall prevented us, just like the rest of the world, from accessing information from the card issuer's servers. Instead, we provided API endpoints only accessible to the credit card firm, using which they provided us with i) users' baseline data during the signup process, and ii) credit card transactions whenever they had been processed. To uphold consumers' interests in the best way possible, we sought to minimize the exchange of data between our partner firm and us. This included us receiving pseudonymized baseline data only after a user had downloaded and explicitly agreed to participate in the study (first step of the signup process); however, the very same baseline data was required on the goal-setting screen (fourth and final step of the signup process). We were able to solve this issue by asking the credit card firm to compute and send the required data for newly registered users with a sufficiently high frequency (every 10 seconds) so that participants were expected to spend enough time on the intermediate two steps of the signup process (in particular, the initial survey).
- Finally, we developed an initial smaller-scale project with less sensitive data being shared in order to allow all involved stakeholders to review the legal and technical setup, and to demonstrate its feasibility. Moreover, a pre-study of our main project was conducted with employees of our research team and the credit card firm to identify and eliminate usability issues and bugs, as well as to validate assumptions made when designing the study [4].

Study Design and Implementation

Weekly credit card spending volumes exhibit large variance both within subjects over time as well as between subjects. We chose weekly feedback cycles (as opposed to daily or monthly ones), to be able to provide feedback with a minimum delay after purchases were made on one hand, and to be able to observe relatively stable spending behavior on the other. Besides excluding customer segments with particularly high variance (e.g. platinum card users, customers with multiple cards, one of which is often given to another household member), we ran power tests and simulated the experiment several thousand times with fully randomized parameters, using spending distribution data provided by our industry partner, and with effect sizes from comparable studies, to estimate the required sample size. We also formulated our main dependent variable relative to the user-specific baseline to accommodate for vastly different spending patterns between subjects.

Another challenge enumerated in the previous section was the incomplete information problem, i.e. participants could reduce their observed spending by switching to other payment means like cash. We addressed this issue by i) designing a meaningful control group with no lesser incentive to hide payments to reach their weekly goals. In addition, we ii) included items both in the initial and final survey, eliciting information regarding possible behavior changes directly. Finally, iii) we plan to estimate the share of wallet of each participant both prior to and during the study [3] in order to detect individuals who significantly changed their behavior during the study.

Insights from User Behavior

Furthermore, conducting a pre-study with employees of the credit card firm as well as our research team (N=95) was helpful in validating design choices, and also for us to get an idea of effect sizes and other variables. This is true even though the pre-study cohort was not sampled from the same population as the main cohort. For instance, weekly spending volumes, app session durations, categorization behavior (e.g. share of expenses categorized as exceptional), number of app visits etc. were all similar to those of the main study. A notable exception to this was the conversion from email invite to fully activated user (which was, unsurprisingly, 5.4 times as high for the pre-study cohort), as well as user retention. 44.4% of the main study cohort stopped using the app within the first three weeks, compared to only 16.8% of pre-study participants. While these figures appear dramatic, it is actually not uncommon for mobile apps to see the vast majority of their users churn early on. Based on our pre-study alone, however, we could have been led into making overly optimistic assumptions regarding user retention.

During our pre-study, we also asked several people to go through the in-app setup process under our supervision, so we could identify possible issues. While this feedback is anecdotal, many users were shocked when they arrived on the goal setting screen, which displayed their average credit card spending of the previous year aggregated on a weekly basis. Since Swiss credit card customers regularly pay their credit card bills in full on a monthly basis, we expected them to have a rather clear picture of the amount they spent every month, or week, respectively. However, it appears that this simple change in aggregation time frame from months to weeks might make the equivalent information more salient for people.

Conclusion & What's next

Drawing from our experience, we can only encourage researchers to engage in ambitious practice-oriented projects even when certain operational aspects seem challenging, or when the involved data is highly sensitive. However, the sensitivity nature of data needs to be addressed in all aspects of such projects, e.g. by appropriate communication towards all internal and external stakeholders (particularly study participants), by a minimally-invasive IT setup with clear data ownership and responsibility policies. Also, a low-key proof of concept project without any academic ambitions might be a smart strategy to establish a trustful collaboration without the stakes being too high (except for the invested time). In our case, this step noticeably helped in getting support for our main project early on. It also helped eliminate technical issues and find solutions to data privacy concerns, which is reflected in our IT architecture, and data separation and data flow concept. We appreciate the challenges that practice-oriented field studies potentially bring along, but we argue that even in industries that traditionally tend to be rather closed, research collaborations can be worthwhile pursuing. Working with consumers in the field proved to be very insightful for us, and we are currently planning a follow-up study with the same partner firm.

Acknowledgements

We thank our project partner, Visa Card Services SA, and particularly Mr. Tony Weber and Mr. Bedrija Hamza for the unique opportunity to conduct research with existing credit card customers, and for their continued support in designing, developing, and publishing the mobile application.

References

1. Anirban Bose and Jean-François Denis. 2017. *World Payments Report 2017*. Retrieved from <https://www.worldpaymentsreport.com/download>
2. Promothesh Chatterjee and Randall L Rose. 2012. Do payment mechanisms change the way consumers perceive products? *Journal of Consumer Research* 38, 6: 1129–1139.
3. Yuxin Chen and Joel H Steckel. 2012. Modeling credit card share of wallet: Solving the incomplete information problem. *Journal of Marketing Research* 49, 5: 655–669.
4. Johannes Huebner, Klaus Fuchs, and Alexander Ilic. 2018. Preventing Over-Spending: Increasing Salience of Credit Card Payments through Smartphone Interventions. In *Multikonferenz Wirtschaftsinformatik (MKWI)*.
5. Hui-Yi Lo and Nigel Harvey. 2011. Shopping without pain: Compulsive buying and the effects of credit card availability in Europe and the Far East. *Journal of Economic Psychology* 32, 1: 79–92.
6. Jill M Norvilitis, Michelle M Merwin, Timothy M Osberg, Patricia V Roehling, Paul Young, and Michele M Kamas. 2006. Personality factors, money attitudes, financial knowledge, and credit-card debt in college students. *Journal of Applied Social Psychology* 36, 6: 1395–1413.
7. Drazen Prelec and George Loewenstein. 1998. The red and the black: Mental accounting of savings and debt. *Marketing science* 17, 1: 4–28.
8. Drazen Prelec and Duncan Simester. 2001. Always leave home without it: A further investigation of the credit-card effect on willingness to pay. *Marketing Letters* 12, 1: 5–12.
9. Priya Raghubir and Joydeep Srivastava. 2008. Monopoly money: the effect of payment coupling and form on spending behavior. *Journal of Experimental Psychology: Applied* 14, 3: 213.
10. Dilip Soman. 2001. Effects of payment mechanism on spending behavior: The role of rehearsal and immediacy of payments. *Journal of Consumer Research* 27, 4: 460–474.
11. Dilip Soman. 2003. The effect of payment transparency on consumption: Quasi-experiments from the field. *Marketing Letters* 14, 3: 173–183.
12. Dilip Soman and Vivian M W Lam. 2002. The effects of prior spending on future spending decisions: The role of acquisition liabilities and payments. *Marketing Letters* 13, 4: 359–372.
13. Abigail B Sussman and Adam L Alter. 2012. The exception is the rule: Underestimating and overspending on exceptional expenses. *Journal of Consumer Research* 39, 4: 800–814.
14. Amos Tversky and Daniel Kahneman. 1975. Judgment under uncertainty: Heuristics and biases. In *Utility, probability, and human decision making*. Springer, 141–162.