Perceived Precautionary Savings Motives: Evidence from FinTech*

Francesco D'Acunto | Thomas Rauter

Christoph Scheuch | Michael Weber

March 17, 2019

PRELIMINARY – PLEASE DO NOT CITE OR CIRCULATE

Abstract

We study the consumption response to the introduction of a mobile overdraft facility on a FinTech app. Users react to the availability of the overdraft by increasing their consumption spending permanently and reallocating consumption from non-discretionary to discretionary goods and services. For identification, we exploit sharp discontinuities in the size of the overdraft limit based on an income rounding rule the app uses to assign credit limits. In the cross section, we find similar responses for young and old users, users with high and low income volatility, and users with steep and flat income paths. The most liquid users—those with high ratios of deposits to income inflows—drive the consumption spending response. These results are not fully consistent with models of financial constraints, buffer stock models with and without durables, present-bias preferences, or the canonical life-cycle permanent income model. We discuss a new channel, the *perceived* precautionary savings channel, which appears consistent with all our results. Under this channel, households with higher liquid wealth behave as if they faced strong precautionary savings motives even though no observables suggests they should do so.

JEL Codes: D14, E21, E51, G21

Keywords: Household Finance, Consumption, Behavioral Finance

*D'Acunto is at Boston College Carroll School of Management, 140 Commonwealth Avenue, Chestnut Hill, MA 02467. Rauter is at the University of Chicago Booth School of Business, 5807 S Woodlawn Ave, Chicago, IL 60637, USA. Scheuch is at the Vienna Graduate School of Finance and WU (Vienna University of Economics and Business), Department of Finance, Accounting and Statistics, Welthandelsplatz 1, Building D4, 1020 Vienna, Austria. Weber is at the University of Chicago Booth School of Business, 5807 S Woodlawn Ave, Chicago, IL 60637, USA and NBER. Email: dacuntof@bc.edu, thomas.rauter@chicagobooth.edu, christoph.scheuch@vgsf.ac.at, michael.weber@chicagobooth.edu. We thank Fabian Nagel and Tom Kim for excellent research assistance. D'Acunto gratefully acknowledges financial support from the Ewing Marion Kauffman Foundation. Rauter and Weber gratefully acknowledge financial support from the IBM Corporation Faculty Fellowship and the Fama Faculty Fellowship at the University of Chicago Booth School of Business. Scheuch gratefully acknowledges financial support from the Austrian Science Fund (FWF project number DK W 1001-G16).

1 Introduction

Credit growth has been a major and often unforeseen driver of financial crises (Schularick and Taylor (2012); Baron and Xiong (2017); Di Maggio and Kermani (2017)). Growing supply of credit to households in particular predicts financial crises over time and across space (Mian and Sufi (2015); Mian et al. (2017)). Understanding micro-level channels through which household credit affects economic cycles is thus not only crucial to deepen our knowledge about the connections between the supply of credit and the business cycle, but also of vital importance to design effective policies that might reduce the disruptive impact of excessive household debt on the real economy.

In this paper, we introduce a unique FinTech setting in which we observe the extensive margin of credit—initiation of overdraft facilities—to document a new empirical regularity that we label *perceived precautionary savings*. Households that display perceived precautionary savings hold deposits equal to more than three times their monthly income, they increase consumption spending substantially after activating the overdraft facility, and substantially more than other households that start using the facility at the same time. Perceived precautionary savers, though, on average do not use the facility, i.e., they never tap into negative deposits. Consistently, their credit scores are less likely to worsen compared to other households that also start using the facility at the same time, and they pay lower fees relative to others.

The income volatility of perceived precautionary savers does not differ from that of other households, either before or after activation of the overdraft facility. Perceived precautionary savers have similar income levels and income growth paths both before and after activation. These households do not transfer money from other accounts in the months before the activation and display similar observable demographics than other households. In particular, age is similar, which makes differences in life-cycle consumption patterns an unlikely explanation for the larger consumption response of perceived precautionary savers.

We label this channel perceived precautionary savings, because perceived precautionary savers behave as if they had precautionary savings motives before accessing the overdraft facility, which they might interpret as an insurance against negative income or unexpected spending shocks. At the same time, they do not display the characteristics we typically relate to precautionary savings motives, such as high income volatility, growing income paths, early life-cycle income paths, or large past spending shocks. Unobservables that capture higher risks of future spending shocks such as medical conditions are also an unlikely explanation for our results, because the average age in our sample is 34 and about 40% of the users display this behavior.

We argue models of buffer stock consumption (Deaton (1991); Carroll (1997)), models of buffer stock consumption with durables (Guerrieri and Lorenzoni (2017); Aydin (2015)), the traditional life-cycle permanent income model of Friedman (1957), or heterogeneous-agents models with assets of different liquidity (Kaplan et al. (2014)) cannot explain our results in full. Our results are closer to Olafsson and Pagel (2018), who document that users of a FinTech app that hold wealth in deposit accounts and use overdraft facilities at the same time increase consumption spending on paydays. In our setting, we show that individuals with the highest ratio of deposits to income react the most in terms of spending once they activate the overdraft facility, whereas in Olafsson and Pagel (2018) individuals with lower amounts of deposits spend more in reaction to income payments.

The main difference between our setting and earlier settings studied in the literature, in which behaviors like the perceived precautionary savings motive were not detected, is the margin of credit we observe. We focus on the first time at which users are allowed to activate an overdraft facility (extensive margin of credit), whereas earlier research mainly focused on increasing the size of existing overdraft facilities (intensive margin of credit). This difference allows us to compare the spending behavior of users before and after activation and hence before and after accessing this form of insurance against negative states of the world for high deposit-to-income users.

When we consider households with lower ratios of deposits to income, we find that they behave in line with traditional hand-to-mouth consumers (Campbell and Mankiw (1989)). After accessing the overdraft facility, they increase their consumption spending slightly, tap into negative deposits, pay higher interests and fees, and are more likely to face worsening credit scores over time.

The European-based FinTech provider with which we cooperate operates in several countries and has more than 1 million customers. This provider introduced a mobile overdraft facility to users who had an online checking account with the bank. The size of the overdraft facility is on average twice the size of users' monthly income. Access to the overdraft facility is logistically simple – users check their eligibility through the online application on their phone and, if eligible, they can activate the facility by clicking on a button. We observe all the transactions users make using their bank card and the type of vendors with which they interact, as well as their credit rating at the monthly level. This setting allows us to assess the effects of providing consumers with a mobile and easy-to-use overdraft facility on their spending habits and credit health both within consumers over time and across consumers whose mobile overdraft has been activated or not.

We propose two empirical designs. First, we estimate a set of double-differences specifications in which we compare the outcomes of app users after activating the overdraft facility relative to before and relative to users whose overdraft facilities are not yet activated. This design focuses on the spending response to the extensive margin of credit, i.e., the provision of a new overdraft facility.

Because the decision to activate the overdraft facility is endogenous to several potential unobservables at the individual level,[1] we propose a sharp regression-discontinuity (RD) design that exploits the rule with which the FinTech provider computes the maximum limit of the overdraft facility. This limit is a rounded function of users' income based on a set of pre-specified income bandwidths of which users are not aware. The FinTech application computes the limits automatically, with no role for bank officers to change the approved limits—this is why the FinTech application allows users to open the approved line of credit automatically in a few seconds—and hence there is no scope for manipulating one's assignment to the high- or low-treatment group in this sharp treatment design. In this RD design, we compare users that are observationally indistinguishable but end up being assigned different overdraft limits based on small differences in their income inflows. This design studies users' spending response to the intensive margin of credit.

We first analyze the panel dimension of our data and discuss average effects in the sample and then discuss the cross-sectional results. We first document that, on average, users that activate the mobile overdraft facility increase their monthly share of consumption spending

[1]Note that we do not detect any systematic differences in the trends of several observables as well as our outcome variables of interest for the period before activation across users that activate or do not activate the mobile overdraft facility though.

over income by 8% at impact and this effect is stable over the months after activation. About three quarters of this increase is due to card-based spending transactions, whereas increased cash withdrawals account for the remaining quarter.

A second fact relates the extent of drawing down from the overdraft facility to the change in consumption spending. We find that 58% of users increase their consumption spending but only 48% ultimately draw down from the facility at least for one month—they ever tap negative deposits and hence borrow from the bank and are responsible to pay interest fees. Moreover, a higher overdraft limit is associated with a higher probability of activating and using the overdraft facility not only for users that ultimately draw down from it, but also for users that never tap into negative deposits.

Third, activation of the overdraft facility does not improve users' credit health in the medium run. On the one hand, one might think the availability of a line of credit might help users build up their credit record and ultimately improve their credit score. On the other hand, the availability of the credit line might worsen credit scores if users roll over their debt and cannot repay the loans they took to finance consumption spending. Consistent with the latter conjecture, we find the share of interest payments and bank fees over users' monthly income increases steadily over time and users' credit scores worsen, on average, without any sign of reversal even several months after activation.

The fourth set of results relates to the change in the composition of users' consumption spending. We find the share of discretionary spending over non-discretionary spending increases substantially with activation. Users increase their shopping and travel expenses on impact, whereas other discretionary expenses such as restaurants and entertainment do not vary substantially. Moreover, the probability users engage in purchasing large-ticket items – which we proxy by transaction sizes larger than 1,000 Euros or 2,000 Euros – increases at impact by about 2 percentage points in the month of activation and stays higher in the subsequent months relative to before activation. Users start to shop for large-ticket items—possibly electronics, cars, or other durable goods—after the facility is made available to them.

The sharp RD design confirms this set of baseline facts alleviating concerns that unobservable determinants of both overdraft activation and changes in consumption drive these results. By construction, the intensive margin treatment we exploit in this analysis is small and we fail to reject the null the spending responses are not permanent over the months after activation. That the effects are not permanent once we eliminate most cross-sectional variation across users suggests that heterogeneity must exist in the consumption response to the overdraft activation.

Indeed, the average results we discussed so far mask large cross-sectional heterogeneity. Whereas the effects are similar across demographic characteristics such as age or gender, Figure 1 documents a striking difference of the effects when comparing users based on the share of their bank deposits over monthly income before activation. The magnitude of the spending response to overdraft activation is indistinguishable from zero for users with the lowest share of deposits to income inflows—those users one might argue are more liquidity constrained—and it increases monotonically with the share of deposits to income, being largest for users with the highest share of deposits over inflows. Users with the highest liquidity in their bank accounts increase their consumption spending scaled by income by about 10 percentage points. As we show below, differences in deposited amounts instead of differences in income levels drive the variation in deposits over income inflows across users.

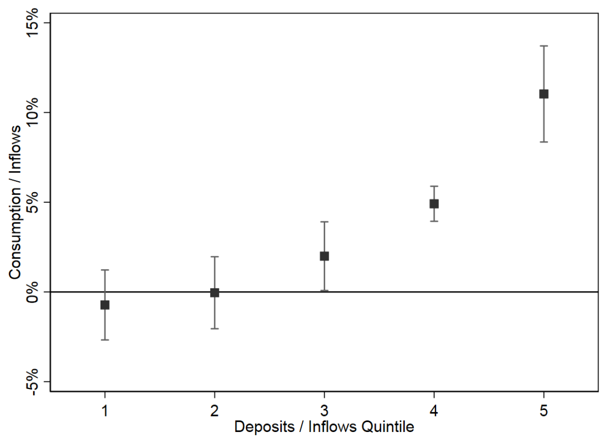

Figure 1:
Spending Response to Overdraft Activation by Deposits to Income Ratio
This figure illustrates the cross-sectional heterogeneity in users' consumption response to the mobile overdraft. To generate the plot, we take the cross-section of users at their treatment date and assign them into non-overlapping quintiles quintiles of deposits to income from lowest to highest. We then interact the resulting grouping variable with a binary indicator that equals 1 if a user has access to a mobile overdraft in given month. Vertical bands represent 95% confidence intervals for the point estimates in each quintile. We double cluster standard errors at the NUTS2 and year-month level.

We discuss a set of possible interpretations for our results. First, that the mobile overdraft facility loosens potential financial constraints on the side of users is implausible, because the users with the highest fraction of cash deposits over income inflows react the most. If anything, the most liquidity-constrained users do not react at all.

Second, access to the overdraft facility might help users smooth consumption in case of growing income paths. Our baseline results are similar across the age distribution and life-cycle consumption patterns are unlikely to matter for older users.[2] To directly test for this mechanism, we also study the income profiles around the overdraft activation and do not find any evidence of different income growth across bins of deposits to income.

Third, the facility might free up liquid resources users were keeping in their bank accounts due to precautionary motives and potential unexpected future income shocks. In this vein, models of buffer-stock savings that allow for both impatience and precautionary savings motives might explain at least in part our results. Our heterogeneity results do not seem fully consistent with this interpretation for a set of reasons. Income uncertainty decreases with age and the heterogeneity results by age we discuss above are not consistent with this form of precautionary savings motive. Moreover, we find the pre-activation volatility of income flows does not predict reaction to the availability of credit. Also, as discussed above, users close to the liquidity constraint do not react, whereas those farther away from the constraint react the most. The buffer-stock interpretation predicts the opposite pattern of reactions.

A fourth potential explanation is a buffer-stock model with durable consumption similar to the one Aydin (2015) studies. In line with this model, we find users mainly spend on durable goods and large-ticket items after activation of the overdraft facility. At the same time, a set of results suggest this interpretation cannot fully explain the results in our setting. In addition to the facts that the least liquidity-constrained users react most and we do not see any differential reaction based on pre-activation income volatility, we find users that react the most on average do not tap into negative deposits and hence *de facto* never use the facility.

A fifth interpretation we consider is present-biased preferences. Users might discount the distant future by more than they discount the immediate future. This interpretation by itself is unlikely to explain all our results, because if the individuals that react the most to overdraft activation were present-biased, they would have consumed out of their higher deposits even *before* activating the overdraft facility, which we do not observe in the data.

We conclude by describing the perceived precautionary savings channel we introduced above. This channel states certain users perceive the need to save for precautionary rea-

[2]Note that our sample is younger than the average population, yet we do detect an age range of 25 years between the bottom and top quintiles of the distribution by age.

son even if objectively they do not face any more uncertain income flows than other users. Both incorrect beliefs about income uncertainty and/or heterogeneity in risk preferences might explain this perception.

Sorting users based on deposits to income ratios can be interpreted as a direct proxy for the extent of the perceived precautionary savings motives because income does barely vary by quintiles. The ratio is thus a proxy for the extent to which users perceive the need of precautionary savings, irrespective of objective measures of income volatility that would not justify higher precautionary savings than other users. In this setting, providing users with insurance in case of possible future negative income shocks reduces precautionary savings motives.

The perceived precautionary savings channel is consistent with the baseline facts we document as well as with the fact that users at the top of the distribution by deposits over income flows react the most to the activation of the overdraft facility. This is because such users restrained their consumption spending believing they would need to hoard precautionary savings and started consuming more after they obtained insurance for future bad states of the world.

The perceived precautionary savings channel predicts a set of additional facts we should be able to observe in our data. First of all, we should see that most of the users with high deposits-to-income ratios increase their consumption but do not tap into negative deposits after activation, so if anything they barely use the insurance the overdraft facility provides to them. Consistently, only 10% of these users have negative deposits but 80% increase their consumption spending in the three months after activation. Instead, about 67% of the users in the bottom of the distribution tap into negative deposits and more than 30% of users in the mid of the distribution do so. Moreover, the average amount of fees—which include both interest amounts paid and transaction fees users pay to withdraw cash—increase by 30% less for the high deposit-to-income ratio users than for the others. Also, even though average credit scores are almost indistinguishable across deposit-over-income inflows bins, users at the top of the distribution are substantially less likely to experience a downgrade (0.03%) relative to users in other bins (about 1.1%).

A potential concern with our interpretation is users might have decided they wanted to purchase big ticket items, moved cash to the deposit account at our bank, and activated the overdraft facility at the same time. But in the data we do not observe users increased their average deposits at our bank in the 3 months before activation.

One might wonder whether our results, and especially the detection of individuals with perceived precautionary savings motives, is peculiar to our sample or whether this phenomenon might be more pervasive. For instance, the majority of the users in our sample reside in Germany and Germans are often labeled as peculiar in their aversion to debt relative to consumers from other countries. In fact, we do not detect any differences in the prevalence of perceived precautionary savings motives across Germans and users in our sample that reside in different European countries. Further research in comparable settings to ours that looks at consumers from other countries is though needed in order to validate the phenomenon we describe in this paper.

Our results also have implication for our understanding of the effects of providing credit to households on business cycle outcomes and the likelihood of financial crises. Our findings suggest that providing insurance against potential negative future spending shocks increases the spending of households with high deposit-to-income ratios while at the same time these households do not take on debt or accumulate interest payments through drawing down from their overdraft facilities. If anything, providing insurance to these households at times of economic slump might increase aggregate demand earlier than would otherwise happen, because those hoarding cash due to perceived precautionary savings motives—who arguably have no objective reasons to not spend and whose increased demand would be important to jump start the economy—would end up spending this cash.

Providing insurance to perceived precautionary savers would be virtually costless based on our results, because these individuals would not end up drawing down from their overdraft facilities. It would also barely resemble the household-credit-led boom and bust cycles we discussed at the beginning, because perceived precautionary savers would not accumulate debt, interest payments, or become riskier borrowers over time. The main challenge to policies based on our findings might be political in nature. At times of economic crisis, economic institutions should provide virtually costless insurance to *wealthier* households (in terms of liquid wealth) to nudge them to spend more instead of providing costly subsidies to poorer households, who might become risky borrowers over time. Whether institutions might find enough political support in parliaments and among voters for this type of policy is a matter of future research inquiry for scholars in political science and sociology.

2 Institutional Setting

We cooperate with a leading European FinTech bank to test for the effects of introducing a mobile overdraft facility on consumption spending behavior. The digital-only bank does not operate a branch network and provides all its services through an Android or iOS mobile app. The bank currently operates under a European banking license in several countries and has more than 1 million customers. Users can open a bank account within 10 minutes by entering their personal information into the app. They are required to verify their identity by providing a copy of their passport or personal ID through video conferencing before the bank confirms the account and users obtain their debit card by mail. The free mobile checking account is the bank's baseline product. The bank does not offer credit cards. Customers manage their account entirely via the bank's mobile app, which provides monthly consumption statistics and allows users to set their daily payment and withdrawal limits, lock their card, or change their pin in real time.

The bank also offers a mobile credit line in several European countries. Residents of these countries with a sufficiently high credit score are eligible for the mobile overdraft facility. Customers activate the credit line directly in their mobile app within one minute and receive a maximum overdraft amount between 500 and 5,000 Euros depending on their credit score and other financial and personal characteristics. The bank uses a fully automated algorithm to allocate maximum credit amounts to users. In Section 5, we describe the bank's loan granting and credit allocation process in detail. Users that are granted a mobile credit line specify their desired credit amount, which they can change in real time via the mobile app depending on their consumption needs. However, customers cannot select an amount that exceeds the maximum overdraft limit allocated by the bank. Users pay an annual interest rate of approximately 10 percentage points on their used overdraft amount, which the bank charges once every calendar quarter. The mobile app provides daily updated information on users' accrued interest costs. Customers can turn on push notifications that remind them whenever their account balance turns negative and they start using the overdraft. The bank cancels the mobile credit line if users default on their interest payments or receive unemployment benefits and experience direct debit reversals.

3 Data and Descriptive Statistics

3.1 Data Sources and Sample Selection

We obtain detailed consumption data, credit line information, and personal user characteristics from a major European FinTech bank. The dataset covers the time period from February 2015 to September 2018. We drop users that do not have access to the mobile overdraft facility because the bank does not offer a credit line in the given European country. We also exclude users that do not satisfy all overdraft eligibility criteria or never applied for the credit line. We focus on individuals that the bank classifies as "main account users" based on their consumption and inflow history to alleviate the concern that customers might have additional accounts with other banks, which we cannot observe. Main account users are individuals that receive a regular monthly salary or incoming standing order into their mobile checking account for at least two consecutive months. Prior research shows that European bank clients satisfy approximately 70% of their daily consumption needs through their primary salary account and that the majority of individuals only have one main account (Bain, 2017; ING, 2018). As a result, our consumption and overdraft data cover most if not all financial activities that main account users carry out via their mobile bank account.

We obtain information about the type, amount, and timestamp of all financial transactions that pass through users' checking account. To protect the identity of its customers, the bank rounded all transaction amounts to the nearest Euro and only provided us with the day but not the exact time of each transaction. The financial transactions covered by our dataset can be classified into 6 broad categories: (i) cash deposits and withdrawals, (ii) incoming or outgoing wire transfers within the Single Euro Payments Area (SEPA), (iii) foreign wire transfers from or to non-SEPA countries, (iv) direct debit withdrawals (including reversals), (v) bank-imposed fees, and (vi) card-based electronic payments. The bank categorizes each electronic payment that a user makes with her debit card into one of seventeen merchant category code (MCC) groups. MCC groups specify the merchant's industry and allow us to identify which type of good or service the account holder purchased. The seventeen MCC groups cover the full range of users' consumption behavior and include both discretionary (e.g., entertainment, shopping, or gastronomy) and non-discretionary consumption types (e.g., groceries, family, or utilities/furniture). Our raw dataset contains 22,419,985 individual financial transactions, which we aggregate into user-month observations. Each within-user time series starts with the month in which the user signed up on the mobile app, verified her identify, and the bank then opened the account, and ends with the closing of the account or the last month of our sample period (September 2018). We code observations of our flow variables as zero if the user did not have any corresponding financial transaction in the given month.

The credit line dataset contains granular information about the application date, granted overdraft amount, and financial characteristics of users that activated the mobile overdraft facility. We observe all user-specific input parameters that enter the bank's credit allocation algorithm, including each individual's credit score, employment status, regular salary, and other credit-relevant inflows. Since the bank shared the precise inner workings of its overdraft granting process with us, we are able to perfectly replicate the credit allocation decision for all mobile credit line users in our sample. Moreover, the credit dataset contains the complete history of all overdraft setting changes that users made once they activated the credit facility. We observe any changes in the actual overdraft usage amount and whether an individual activated push notifications that pop-up whenever the account balance turns negative and users start drawing on the credit line.

The bank also provides us with demographic and personal information about each main account user. We obtain data on users' gender, year of birth, and residential zip code. To ensure data anonymity, the bank does not share the name, address, or precise date of birth of its account holders with us. In Appendix A, we define all variables that we use in our empirical analysis.

3.2 Descriptive Statistics

Table 1 provides descriptive statistics for our main overdraft sample. We trim all ratios that involve consumption-related variables at the 5th and 95th percentile to mitigate the impact of outliers due to data errors or extreme values.[3] Our dataset contains 595,244 user-month observations of 40,979 individuals who obtained a mobile credit line between February 2015 and October 2017. The user base of the FinTech bank consists primarily of male Millennials who live in urban areas. The average user in our sample is 34 years old, has monthly inflows of 2,073 Euros, and opened the mobile bank account 1.3 years ago. 80% of our sample users are male and 52% live in large cities with more than 500,000 inhabitants. All overdraft users combined spent a total of 435 million Euros via their mobile checking account over our sample period. On average, these individuals consume 54% of their (lagged) monthly inflows, of which approximately two thirds are attributable to electronic card transactions and the remainder is cash consumption. For each Euro that users spend electronically on non-discretionary goods, they purchase 81 cents of discretionary items.

Main account users have access to the bank's mobile overdraft facility in 88% of all user-months. However, these individuals only actively use the credit line 50% of the time. The average maximum overdraft amount equals about 1,143 Euros, of which the clients in our sample typically use 62%. The majority of active credit line users allow push notifications on their mobile phone that notify them whenever their account balance turns negative and they start accruing overdraft interest costs.

4 The Effect of Mobile Overdrafts on Users' Consumption Behavior and Credit Risk

In this section, we examine how main account users change their consumption behavior once they activate an overdraft facility on their mobile phone. Specifically, we provide evidence on the effect of mobile credit lines on users' overall level of consumption and examine whether mobile overdrafts are associated with higher or lower credit quality among those individuals that actively use them.

[3] We obtain similar results when we instead winsorize our regression variables at the 5th and 95th percentile or when we use alternative trimming approaches (e.g., trimming at the 1% level in each tail).

4.1 Overall Consumption Response

We begin our empirical analysis by examining users' monthly consumption expenditures and use a difference-in-differences (DD) design to identify the effect of mobile credit facilities. The DD estimator compares changes in the level of consumption around the activation of mobile overdrafts between individuals that did or did not yet use the credit line. We estimate our treatment effects within mobile overdraft users to condition the analysis on the endogenous selection into the credit facility and to alleviate the concern that individuals who never activate the overdraft might be so fundamentally different from treated users that their consumption patterns in the pre-overdraft period are not parallel to each other. Specifically, we estimate the following baseline OLS regression model:

$$\text{Consumption}_{i,t} = \beta \times \text{Overdraft Available}_{i,t} + \text{Fixed Effects}_{i,t} + \varepsilon_{i,t} \tag{1}$$

The dependent variable is the sum of all cash withdrawals and card-based purchase transactions by individual i in month t, divided by the amount of the user's account inflows in month $t-1$. Our main variable of interest *Overdraft Available* is an indicator variable equal to one beginning in the month in which the account holder got access to the credit facility on her mobile phone.[4]

We include user fixed effects to control for time-invariant variation in consumption patterns between overdraft users resulting from differences in occupation, gender, cultural backgrounds, or education. We add NUTS3×year-month fixed effects to account for concurrent but unrelated time-varying economic or institutional changes within local, sub-national districts (that contain a maximum of 800,000 inhabitants). We double cluster standard errors at the NUTS2 and year-month level since consumption patterns are likely correlated cross-sectionally and over time within a given administrative district.

In Table 2, we present regression results for the estimated effect of mobile overdraft availability on consumption behavior. In column (1), we report the results of our baseline specification. The estimated consumption effect of mobile overdrafts is positive and highly statistically

[4] In Table IA1 of the Internet Appendix, we instead define our treatment indicator based on whether individuals actually use the mobile credit line (*Overdraft Enabled*) and find both qualitatively and quantitatively similar results. We use the *Overdraft Available* dummy as our baseline treatment indicator since users might already adjust their consumption behavior when they obtain access to the credit line, realizing that they can enable the overdraft on their mobile phone anytime they want going forward.

significant. The coefficient magnitude indicates an increase in overall cash and card-based consumption of 4.413 percentage points (t-statistic: 11.30) of the user's lagged account inflows, which corresponds to an increase of approximately 8.12% relative to the sample mean (4.413/54.36). In columns (2) and (3), we differentiate between cash- and card-based consumption and find that account holders spend significantly more through both payment types. In relative terms, the increase in card-based consumption (coefficient: 2.952; t-statistic: 8.83) is larger than the increase in cash withdrawals (coefficient: 1.051; t-statistic: 6.29) and accounts for approximately 75% of user's overall consumption response.

In Figures 2 and 3, we provide graphical evidence that overdraft and control users have parallel and almost identical consumption patterns during the time period leading up to the mobile overdraft activation. Individuals sharply increase their spending during the first two months in which they can access the credit line. After that, users' consumption reverts to a stable treatment effect of approximately 2 percentage points of monthly account inflows.

4.2 Credit Quality Implications for Overdraft Users

Mobile overdrafts also likely affect the credit risk of individuals that activate the credit line. Conceptually, it is unclear whether these credit lines hurt or improve users' credit score. On the one hand, mobile overdrafts by FinTech lenders might serve as a gateway to credit for those individuals that would have otherwise not received a credit line from traditional brick and mortar banks through face-to-face lending (e.g., Bartlett et al., 2018; Claessens et al., 2018; Dobbie et al., 2018). In this context, FinTech-based consumer lending could facilitate financial inclusion and help previously underbanked individuals to build credit. On the other hand, mobile overdrafts could induce users that already had access to consumer credit to exceed their debt capacity, default on their payment obligations, and thereby worsen their credit score (e.g., Ausubel, 1991; Laibson et al., 2007; Stango and Zinman, 2009).

To assess the credit quality implications of mobile credit lines, we compare changes in users' credit score for main account holders that activated the overdraft and individuals that did not yet (but eventually will) apply for the credit facility:

$$\text{Credit Rating}_{i,t} = \beta \times \text{Overdraft Available}_{i,t} + \text{Fixed Effects}_{i,t} + \varepsilon_{i,t} \tag{2}$$

Credit Rating is the consumer credit rating of user i in month t. We only consider individuals with a minimum rating of "F" who are thus eligible for the overdraft. We map categorical ratings into natural numbers using a scale from 1 to 6, where 1 corresponds to the highest rating ("A") and 6 to the lowest rating ("F"). We include the same fixed effects as in equation (1). The bank requests users' credit ratings from consumer credit bureaus at the time of the overdraft application and continuously tracks individuals' credit history once they have been granted a credit line. To minimize costs, the FinTech lender does not request and monitor ratings before users apply for the credit line. Since we obtain ratings information from our collaborating bank (and not credit bureaus that quantify individuals' credit quality on an ongoing basis both before and after an overdraft application), we are only able to examine the evolution of individuals' credit rating relative to the time when the account holder applied for the mobile credit line, but not any time before that.

In Table 3, we present the regression results for the estimated effect of mobile overdraft facilities on users' credit rating. In column (1), we estimate equation (2) for all main account users who eventually receive a credit line. We find that the coefficient of *Overdraft Available* is positive and statistically significant (t-statistic: 2.81), indicating that the credit risk of users increases once they activate the overdraft. The adjusted R-squared of 0.998 suggests that individual credit ratings are stale and that our user fixed effect therefore explains almost all of the variation in users' credit risk. Indeed, only 63 main account users with ratings data experience changes in their credit rating over our sample period. To mitigate the concern that our regression model might overfit the ratings data and thereby bias our inference, we reestimate equation (2) only for users whose credit rating changes at least once following the overdraft application. In column (2), we find that our overdraft coefficient remains positive and becomes even more significant (t-static: 3.21). In economic terms, the increase of 0.99 rating notches suggests that users' credit risk deteriorates by approximately 24% relative to the sample mean (0.992/4.175). Figure 4 shows that decrease in credit quality occurs soon after the activation of the mobile credit line. These treatment dynamics alleviate the concern that other confounding factors drive the results since remaining omitted variables would need to be correlated with individuals' credit ratings and the entire distribution of overdraft activation dates across all main account holders.

4.3 Overdraft Usage and Fees

Having established the baseline results, we next validate our research design and examine whether individuals actually use the credit line once they obtain access to a mobile overdraft. We estimate the following OLS regression model:

$$\text{Overdraft Usage}_{i,t} = \beta \times \text{Overdraft Available or Amount}_{i,t} + \text{Fixed Effects}_{i,t} + \varepsilon_{i,t} \quad (3)$$

The dependent variable *Overdraft Usage* is either: *Overdraft Enabled* or *Negative Deposits*. *Overdraft Enabled* is a binary indicator that takes the value of one during months in which the user activates the credit line conditional on having been granted an overdraft. *Negative Deposits* is a dummy variable that equals one whenever the user has a negative account balance and therefore draws on the credit line. *Overdraft Amount* is the natural logarithm of the user's maximum overdraft amount. While *Overdraft Available* (our main treatment indicator) captures access to credit at the extensive margin, *Overdraft Amount* quantifies credit availability at the intensive margin. The fixed effects mirror those of equations (1) and (2).

In Table 4, we find that individuals who obtain access to a mobile overdraft do in fact use the credit line. The regression coefficients in columns (1) and (3) indicate that, conditional on being granted a credit line, the probabilities of users activating the overdraft and tapping into negative deposits equal 0.708 (t-statistic: 28.07) and 0.482 (t-stat: 29.26), respectively. At the intensive margin, our results suggest that users with higher maximum overdraft limits are more likely to use the credit facility. In economic terms, the estimate in column (2) (column (4)) suggests that a one-standard deviation increase in *Overdraft Amount* is associated with a 30.72% (54.24%) higher likelihood that users enable the mobile overdraft (face a negative account balance) relative to the sample mean (i.e., 0.154/0.501 and 0.175/0.322).

Main account holders accrue an annual interest rate of approximately 10% percentage points whenever they draw on the credit line. The bank computes individuals' overdraft usage times on a quarterly basis and charges overdraft fees at the beginning of each subsequent calendar quarter. To examine the effect of mobile credit lines on users' account fees, we thus estimate the following regression model:

$$\text{Fees}_{i,q} = \beta \times \text{Overdraft Available or Enabled or Amount}_{i,q-1} + \text{Fixed Effects}_{i,q} + \varepsilon_{i,q} \quad (4)$$

The dependent variable is the total amount of fees that the bank deducts from the checking account of individual i in quarter q, normalized by the user's contemporaneous account inflows. The definitions of our *Overdraft Available*, *Overdraft Enabled*, and *Overdraft Amount* variables are identical to those of the previous analyses. We add the same fixed effects as in equations (1) to (3). Since our outcome variable only varies once per calendar quarter, we collapse our data to the user-quarter level. We adjust standard errors for within group clusters at the level of users' NUTS2 region of residence but do not cluster by calendar quarter since the low number of only 15 quarters could result in an overrejection of the null hypothesis (Cameron and Miller, 2015).

In Table 5, we report the regression results for the estimated effect of mobile overdrafts on users' quarterly account fees. Consistent with the bank's terms and conditions, we find that users pay higher overall fees following the activation of the credit line. We document a strong positive association between individuals' overdraft usage and their normalized account fees, both at the extensive and the intensive margin. The point estimate of 0.353 (t-statistic: 4.07) in column (1) indicates an increase in users' account fees of about 34% relative to the sample mean (i.e., 0.353/1.041) once the overdraft becomes available. In Figure 5, we provide graphical evidence that users with access to a mobile credit line and those without have similar fee patterns during the pre-overdraft period. Moreover, Figure 5 shows that the account fees of overdraft users increase sharply soon after the bank starts charging overdraft interest costs.

Overall, the results in this section indicate that users who apply for and obtain access to a mobile overdraft also draw on and pay fees for the credit line. Conditional on activating the overdraft, users substantially increase their monthly cash- and card-based consumption expenditures, particularly during the first 2 to 3 months. Finally, our credit risk analysis shows that consumer credit ratings decline once individuals start using the mobile overdraft.

5 Regression Discontinuity Analysis

So far, our findings indicate that access to mobile overdrafts is associated with higher overall consumption by those individuals that apply for and obtain a credit line. However, these results do not speak to whether mobile overdrafts have a causal impact on users' consumption behavior. Assessing the causal effects of mobile credit facilities entails several identification challenges. First, causality often runs in both directions as users activate credit lines in anticipation of higher future consumption (reverse causality). Second, correlated omitted variables might simultaneously impact users' consumption behavior and overdraft activation decision, giving rise to a spurious relation between the two. One example for such a correlated omitted variable might be time-varying, user-specific exposure to television commercials that independently advertise the bank's overdraft and various consumer products. In this section, we address these endogeneity concerns and estimate the causal effects of mobile overdrafts in a sharp regression discontinuity (RD) design that exploits variation in users' overdraft limits based on thresholds embedded in the bank's credit allocation algorithm. Our sharp RD design conditions the analysis on users' (endogenous) selection into the mobile overdraft and relies on exogenous variation in the size of the credit line along the intensive margin.

5.1 Credit Allocation Algorithm

The bank's credit allocation process consists of two steps. First, the bank determines whether users pass all exclusion criteria and are thus eligible for a mobile credit line. Overdraft applicants receive a credit line if they (i) are employed, (ii) live in countries where the bank offers a mobile overdraft, (iii) have a minimum credit score of F, and (iv) their checking account did not trigger any direct debit reversals. The bank obtains credit scores from consumer credit bureaus, which collect information on users' credit histories to estimate default probabilities and assign individual credit ratings from A (lowest default risk) to M (highest default risk). A credit score of F implies that the individual has an estimated default probability of less than 10 percent.

Second, the bank determines the maximum overdraft amount for each eligible user based on the applicant's credit score and average account income according to the following formula:

$$\text{Overdraft Amount} = \begin{cases} \text{Max Limit} & \text{if } 2 \times \text{Income} \geq \text{Max Limit} \\ \text{Min Limit} & \text{if } 2 \times \text{Income} \leq \text{Min Limit} \\ 250 \times \lfloor \frac{2 \times \cdot \text{Income}}{250} \rceil & \text{otherwise} \end{cases} \tag{5}$$

where $\lfloor x \rceil$ rounds the number x to the nearest integer.

For each rating notch between A and F, the bank specifies a lower (*Min Limit*) and upper limit (*Max Limit*) for each user's allocated credit amount. Income is a linear function of the user's different inflow types in the months prior to the overdraft application. Our data sharing agreement with the bank does not allow us to report the rating-specific overdraft limits or the precise formula that transforms users' account inflows into Income. However, we can disclose that the bank differentiates between regular salary and non-salary related inflows (e.g., pension, child benefits, study support from parents etc.) and puts a higher weight on the former. The lower and upper overdraft limits monotonically increase in the customer's credit rating and range between 500 and 5,000 Euros.

To determine each user's maximum available overdraft amount, the bank's fully automated credit allocation algorithm multiplies the Income variable by 2. If the resulting value exceeds (falls below) the upper (lower) credit limit, the maximum overdraft amount is bounded from above (below) by the rating-specific limit. If the doubled Income falls in between the upper and lower limit, the amount is rounded to the closest 250 Euro multiple at the midpoint. Panel A of Figure 6 illustrates the rounding convention embedded in the credit algorithm. For example, if overdraft applicant A has a salary of 2,100 Euros and no additional account inflows, her implied overdraft amount after multiplying the income by 2 equals 4,200 Euros, which, if rounded to the closest 250 Euro threshold, translates into a maximum available overdraft amount of 4,250 Euros (assuming that the upper and lower credit limits do not bite). The bank's credit allocation process gives rise to 18 unique thresholds in the interval between 500 and 5,000 Euros, at which the maximum overdraft amount jumps discontinuously by 250 Euros. At these thresholds, users with almost identical Income that find themselves on opposite sides of the rounding threshold receive different overdraft limits for plausibly exogenous reasons.

5.2 Empirical Implementation

We limit our analysis to users whose maximum overdraft amount equals the individual's income multiplied by two and rounded to the nearest multiple of 250. That is, we drop all users whose transformed income exceeds or falls below the upper or lower credit limit (within the given rating notch) such that the rounding thresholds embedded in the bank's credit allocation algorithm do not affect the maximum overdraft amount. For each user in our RD sample, we compute the forcing variable X_i, which quantifies the individual's distance (in Euros) to the nearest rounding threshold. X_i removes differences in absolute rounding thresholds across individuals and is centered around zero. Users with $X_i \geq 0$ are treated and receive a maximum overdraft amount that is 250 Euros higher than those of control users for whom $X_i < 0$. The probability that a user's overdraft limit gets rounded up by 250 Euros changes discontinuously from 0 to 1 at the rounding threshold. Panel B of Figure 6 illustrates the exact treatment rule of our sharp RD design and plots users' treatment assignment for different values of the forcing variable X_i. In areas close to the rounding threshold (where $X_i = 0$), treated and control users have almost identical income profiles.

To examine the causal effect of mobile credit lines on users' consumption behavior, we implement the following sharp RD design:

$$\tau \equiv \mathbb{E}\left(C_i(1)|X_i = 0\right) - \mathbb{E}\left(C_i(0)|X_i = 0\right). \tag{6}$$

τ is the RD treatment effect and $C_i(1/0)$ is the change in treated (1) or control (0) user's average consumption 3 months before and after the credit allocation decision, divided by the individual's average inflows in the 3 months prior to the overdraft application. To estimate this model, we fit a weighted least squares regression of the observed consumption change on a constant and polynomials of X_i on both sides of the rounding threshold. The RD treatment effect is the difference in estimated intercepts from these 2 local weighted regressions. Formally, each user's consumption change equals:

$$C_i = \begin{cases} C_i(0) & \text{if } X_i \geq 0 \\ C_i(1) & \text{if } X_i < 0. \end{cases} \tag{7}$$

We focus on observations within the interval $[-h, h]$ around the rounding threshold, where $h > 0$ denotes our bandwidth of choice. The kernel function $K(\cdot)$ specifies our regression weights. $\hat{\mu}_{+/-}$ is the estimate of $\mathbb{E}\left(C_i(1/0)|X_i = 0\right)$ for observations above or below the threshold, which we define as:

$$\hat{C}_i = \hat{\mu}_{+/-} + \sum_{j=1}^{p} \hat{\mu}_{+/-,j} X_i^j, \tag{8}$$

where p denotes to the order of our local polynomial. The RD treatment effect then equals:

$$\hat{\tau} = \hat{\mu}_{+} - \hat{\mu}_{-}. \tag{9}$$

To operationalize the RD estimator, we need to specify (i) the order of polynomial p, (ii) the kernel function $K(\cdot)$, and (iii) the bandwidth h. We follow Gelman and Imbens (2018) and only use polynomials of order 1 and 2 to avoid overfitting issues. We apply weights from a triangular kernel because it is the mean squared error (MSE) minimizing choice for point estimation in our context (Cheng et al., 1997). Finally, we employ the MSE-optimal bandwidth selection procedure recommended by Calonico et al. (2014), which corrects for the non-negligible bias resulting from subjective bandwidth choices. We residualize the outcome variables of our RD analysis with country×year-month fixed effects to ensure that we compare treated and control users from the same European country at a similar point in time.

5.3 Assessing Identification Assumptions: Treatment Manipulation and Balancing Tests

Our sharp RD design critically relies on the assumption that the forcing variable for individuals just below the threshold is similar to those just above the threshold. If users can manipulate the forcing variable and thereby their assignment to treatment and control groups, this local continuity assumption is violated, which results in biased RD estimates (Roberts and Whited, 2013).

Conceptually, it is unlikely that users can control their treatment assignment in our setting. Most importantly, the bank's credit allocation algorithm is proprietary information and not known by overdraft users. Even if individuals were informed about the precise inner workings of the overdraft allocation formula (in particular its rounding thresholds), it seems implausible that users could precisely manipulate their income, for example, by negotiating a higher wage with their employer (Lee and Lemieux, 2010). Moreover, it appears unlikely that overdraft users would be willing to voluntarily forgo parts of their salary just to obtain access to a higher credit limit.

To formally assess the validity of the local continuity assumption, we test for the presence of a discontinuity in the density of X_i at the rounding threshold. If users systematically inflate their income to receive a higher overdraft limit, we should observe a kink in the distribution of our forcing variable right above the threshold. We use the local polynomial density estimator of Cattaneo et al. (2017) to test whether overdraft users manipulate their assignment into treatment and control group. In Figure 7, we plot both the frequency distribution (Panel A) and density function based on quadratic local polynomials (Panel B) of our running variable and do not find graphical evidence for bunching above the rounding threshold. In Table 6 and Figure IA3 in the Internet Appendix, we report the estimation results of the formal treatment manipulation test by Cattaneo et al. (2017) for different polynomial and bandwidth choices. In all specifications, we fail to reject the null hypothesis that our running variable is locally continuous around the rounding threshold.

The local continuity assumption implies that individuals below and above the cutoff should not only be similar in terms of the forcing variable but also along other characteristics. Since overdraft users lack the ability to precisely manipulate their distance to the rounding threshold, there should not be systematic differences in observable characteristics between the two groups of individuals. Consistent with this argument, in Table 7, we do not find significant differences in the average age, gender, time since account opening, rating, account inflows, and consumption between treated and control users prior to the activation of the mobile overdraft. As an alternative balancing test, we repeat our RD analysis but replace our main outcome variable with each observable user characteristic. In Table IA5 of the Internet Appendix, we document that the RD treatment effect for all our covariates is economically and statistically indistinguishable from zero.

Overall, the evidence in this subsection indicates that overdraft users do not manipulate their treatment assignment and that individuals above and below the rounding threshold have similar observable characteristics. Both findings suggest that the local continuity assumption is satisfied and thereby corroborate the internal validity of our sharp RD design.

5.4 RD Consumption Effect on Impact and Over Time

In Figure 8, we graphically illustrate the RD treatment effect of a 250 Euro higher overdraft limit on users' consumption behavior. We aggregate our data into disjoint bins and make sure that each bin contains either treatment or control observations. We then calculate the average value of our outcome variable, plot this value above the midpoint of the bin, and separately fit 2 linear regressions through all observations on each side of the rounding threshold. We choose the number of bins based on the evenly-spaced mimicking variance method by Calonico et al. (2015).

In Panel A of Figure 8, we verify that individuals just above the threshold indeed receive a higher maximum overdraft amount (relative to their income) compared to users just below the threshold. The slope of both fitted regressions lines is negative since, within treatment and control group, individuals with larger values of our forcing variable X_i have a higher income, which we use to normalize users' overdraft limit. In Panel B, we plot the change in average (normalized) consumption three months before and after the user obtained access to the credit line. We find a positive discontinuity in users' consumption growth right at the rounding threshold, indicating that treated users consume more relative to control users following the exogenous assignment of a 250 Euro higher overdraft limit.[5] In Table 8, we present the coefficients from estimating the sharp RD design we formalized in equations (6) to (9). We estimate first- and second-order polynomial regressions at the rounding threshold and report both bias-corrected and conventional t-statistics (Cattaneo et al., 2017; Gelman and Imbens, 2018). In columns (1) and (2), we document a positive and highly statistically significant RD treatment effect on users' consumption growth of between approximately 11 and 14 percentage points. The coefficient estimates do not attenuate when we add user characteristics as control variables in a linear and additive-separable way. In line with Calonico et al. (2018), we find that adding covariates increases the precision of our point estimates, which again suggests

[5] Our results are both qualitatively and quantitatively similar if we normalize the change in users' consumption by income rather than inflows.

that the local continuity assumption is satisfied. The coefficient magnitude of 10.87 in our most conservative specification implies that treated overdraft users consume approximately 203 Euros of the 250 additional Euros that they receive as credit, which corresponds to a marginal propensity to consume (MPC) of about 81%.[6]

We conduct two robustness tests to assess the sensitivity of our RD estimates. First, we examine how sensitive the RD results are with respect to the choice of our bandwidth (Imbens and Lemieux, 2008). Varying the bandwidth is only meaningful over small intervals around the MSE-optimal choice (Cattaneo et al., 2020). Bandwidths much larger than the MSE-optimal bandwidth bias the RD estimator, while substantially smaller bandwidths inflate its variance. Figure IA4 shows that different bandwidth choices do neither substantially affect the magnitude of the point estimate, nor its significance. Second, we assess how robust our RD point estimates are to excluding data close to the threshold (e.g., Barreca et al., 2011, 2016). We drop users located within the radius $r > 0$ of the rounding cutoff, that is, we exclude observations for which $|X_i| \leq r$ (Cattaneo et al., 2020). Figure IA5 plots the coefficient estimates for different choices of r and shows that observations close to the rounding threshold do not drive our results.

In Figure 9, we map out the RD treatment effect over time. Consistent with the results in Section 4, we find that treated users only increase their consumption during the first three months after they receive the higher overdraft limit. Overall, the findings of our RD analysis indicate that mobile overdraft facilities have a positive causal effect on users' consumption behavior that is temporary.

6 Heterogeneity and Economic Mechanisms

Our results so far have focused on average effects across users. To better understand the economic channels that might drive our results and to assess the extent to which existing consumption models might explain them we turn on to study the heterogeneity of our baseline

[6] We compute the marginal propensity to consume as MPC $= [\hat{\tau} \cdot \text{Inflows}_{t-3:t-1} + \text{Consumption}_{t+1:t+3}(0) - \text{Consumption}_{t-3:t-1}(0)]/250$, where $\hat{\tau}$ is our estimated RD treatment effect, $\text{Inflows}_{t-3:t-1}$ is the average inflows before activation of the overdraft (identical for treated and control users), $\text{Consumption}_{t-3:t-1}$ is the average consumption in the pre-period (identical for treated and control users), and $\text{Consumption}_{t+1:t+3}(0)$ is the average consumption of control users during the three months after the treatment.

results in the cross section of users (Jappelli and Pistaferri (2017)).

6.1 Heterogeneity: Income Growth, Age, Liquidity

We assess the heterogeneity of the baseline effect across two types of splits—2 groups (below and above the median by each characteristics) and 5 groups (across quintiles of the distribution). The cross-sectional variables we consider are the three characteristics we can observe in the data, that is, (i) the growth of income inflows 6 months after overdraft activation relative to 6 months before activation; (ii) users' age; and (iii) the share of deposits to income inflows in the month prior to activation, which capture users' share of liquid resources over monthly income.

Table 10 reports the results in tabular form for estimating our baseline regression and including interaction terms for households above the median for each characteristic.

The lifecycle permanent income hypothesis (LFPIH) suggests that agents want to smooth consumption. Empirically, income paths are increasing early in life before flattening out. Hence, the LFPIH predicts younger users and users with a steeply increasing income path should be more likely to use the overdraft facility to smooth consumption and increase their spending relative to income once they have access to the overdraft facility. In columns (1) and (2), instead, we see no heterogeneity exists by age. Users with above-median income growth around the overdraft activation are, if anything, increasing their consumption spending by less than users with income growth below the median. These two results appear inconsistent with the LFPIH.

Column (3) of Table 10 splits the sample in two groups based on the share of deposits over income inflows in the month before activation to proxy for users' liquid wealth and the binding of financial constraints. Here, we detect substantial heterogeneity. In fact, users whose deposits-to-income ratio is below the median do not appear to change their consumption spending over income at all after activation, neither economically—if anything, the points estimate is slightly negative (-0.096)—nor statistically (t-stat=0.13). Instead, the effect is about twice as large than the average baseline effect in the sample for users whose deposits-to-income ratio is above the median (7.153, t-stat=7.58).

To dig deeper and interpret these heterogeneity results, we move on to split the sample into quintiles based on each of the three characteristics discussed above. We repeat the baseline regressions adding a set of interactions between a dummy for whether the user belongs to each quintile with the dummy variable for observations before and after activation. To make the results easier to visualize, we report the estimated coefficients and 95% confidence intervals in graphical form in Figure 13.

The top panel of Figure 13 reports the effects across quintiles of income inflows growth around activation. Consistent with Table 10, the estimated effect is twice as large for users in the first quintile (4.5%) relative to users in the top 3 quintiles (2%), although we do not reject the null hypothesis that the effects are equal across any of the quintiles, including when we compare the size of the effect in the first and fifth quintiles.

In the middle panel of Figure 13, we split the sample into quintiles by age. Despite the fact that our users are on average younger than the broader population, we still detect substantial differences in age between the bottom quintile and the top quintile, whose averages are about 20 years old and 45 years old. We can plausibly argue users across these quintiles are on different consumption life-cycle paths. The panel shows the non-result in column (2) of Table 10 is not driven by nonlinear heterogeneities of the effect in the age distribution, but instead the effect is stable across the whole distribution. In terms of magnitudes, we estimate coefficients that range between 3% and 4% for the effect in each quintile. Moreover, we do not reject the null that the effects are the same across all quintiles at any plausible level of significance.

The bottom panel of Figure 13 considers quintiles by the deposits-to-income ratio in the month before activation. We detect substantial heterogeneity in this case. Whereas the effect is (insignificantly) negative for users in the bottom quintile—the most liquidity-constrained users—and is zero for those in the second quintile, it increases nonlinearly at higher quintiles. The estimated effect is about 2% for users in the third quintile, 5% in the fourth quintile, and 12% in the top quintile.

6.2 The Perceived Precautionary Savings Mechanism

The heterogeneity results suggest a pattern whereby users with higher liquidity (cash deposits) over income react more than others to the activation of the overdraft facility in terms of consumption response. This pattern is intriguing, because we might have expected the most liquid users were those that had the least need of a overdraft facility if they wanted to spend before activation. To the extent that the overdraft facility is mainly used to smooth spending and loosen liquidity constraints, we might have expected users in the bottom quintile by deposits over income would have reacted the most instead of the pattern we observe in the data.

Users that hold substantial liquidity might change their behavior after they access the overdraft facility due to precautionary savings motives and the need to maintain enough liquidity available in case of potential future negative income shocks. To understand if this motivation can explain our findings, we consider a set of dimensions earlier research has associated with precautionary savings motives. If our results captured traditional precautionary savings motives, we should expect these dimensions vary systematically in the cross section of users by deposits-to-income ratios before activation of the overdraft facility. Table 11 reports the results for these cross-sectional sample splits.

First, we consider two demographic characteristics—age and gender. Age increases monotonically with the bins by deposits-to-income inflows, ranging from 31.5 to 34.9. Although the difference between the fifth and first bin is statistically different from zero, the magnitude of this difference is less that 10% of the average in the top bin. Most importantly, though, the typical precautionary savings explanation would suggest that *younger* users have higher precautionary savings motives, because these users are likely to have more uncertain income flows in the near future, might expect higher income growth paths, might face less employment stability, and might still not be in the workforce at all. As far as gender is concerned, we barely detect any economic differences in the share of women across bins. Although the highest bin includes a share of women (23.2%) that is statistically different from the first bin (19.4%), the magnitudes of these differences are small and cannot explain the substantially different reactions in terms of consumption spending we document in Figure 13.

We move on and dig deeper into the differences across bins. An important fact to establish is whether the variation in the deposits-to-income ratios across bins is mainly driven by differences in the amounts of deposits, as our liquidity interpretation of this ratio suggests, or by differences in income flows. Table 11 documents that variation in the amount of deposits drives the variation. Deposits before activation increase monotonically across bins and are substantially higher in the top bin relative to all other bins, being close to zero in the first bin. Instead, the income inflows before activation vary marginally across bins and do not increase monotonically.

Notably, the average income inflows in the bottom and top bins are almost indistinguishable economically and statistically.

As discussed when considering age profiles, younger users might face if anything higher variation in their income inflows than older users. To understand if instead this conjecture is wrong and the standard deviation of income inflows drive our results, we look at this variable directly in Table 11. Specifically, we average the standard deviation in income inflows over the 12 months before overdraft facility activation across bins. We do not detect any systematic patterns or differences between the bottom and top bin by deposits-to-income ratios. This fact is prima facie evidence that differences in past variation of income flows do not justify why users in the top bin behave as if they have stronger precautionary savings motives.

Another typical driver of precautionary savings motives is a fast income growth profile over time. Even in this case, irrespective of whether we compute income growth for the three months or six months before and after activation, we do not detect any systematic patterns or economically/statistically significant differences across bin by deposits-to-income ratios.

So far, comparing bins by deposits-to-income ratios does not suggest that users in the top bin have any objective reasons to hold stronger precautionary savings motives than users in the lower bins. We move on to assess whether, apart from increasing consumption spending after activating the overdraft facility, users in the top bin also behave in line with precautionary savers in terms of overdraft facility usage. In particular, precautionary savers, contrary to liquidity-constrained individuals, would likely not tap into negative deposits and would not raise debt through the overdraft facility. Instead, they should view the facility as a form of insurance against negative income shocks and would thus spend some of the existing liquidity they had accumulated before the ovedraft facility was available once they know they can tap into negative deposits if needed.

The results in Table 11 are broadly consistent with the users in the top bin by deposits-to-income behaving as if they had strong precautionary-savings motives. First, these users are substantially less likely than users in lower bins to tap into negative deposits after activation, despite increasing their consumption spending relative to the pre-period substantially more than them. The probability of tapping into negative deposits ranges from 67% for users in the bottom bin to 10% for users in the top bin. Consistently, users in the top bin pay less fees, on average, relative to users in the bottom bin. Note the fee variable we observe pools together interest payments for using the overdraft facility and bank fees like late-payment fees or cash withdrawal fees. Finally, users in the top bin are less likely to face credit score downgrades after activating the facility relative to users in lower bins, and the likelihood of downgrading declines monotonically with the deposits-to-income ratio. The probability of downgrades, instead, was not different across bins before activation and was very close to zero, on average, in all bins.

Because the 75^{th} percentile of age in our sample is 38.5, it does not even seem plausible to assume that a large fraction of the individuals in our sample might have objective precautionary savings motives due to potentially unexpected large medical bills or other medical-related expenses. A potential concern with our interpretation is that users decide they want to purchase big ticket items and move cash to the deposit account at our bank before they activate the overdraft facility. But in the data we do not observe users increase their average deposits at our bank in the 3 months before activation.

Overall, users with a high share of deposits to income ratios and hence high liquidity behave as if they had precautionary savings motives and saved before the overdraft facility became available to them. Once they have access to the overdraft facility which acts like an insurance for additional future spending needs, possibly due to unexpected spending needs or income shortfalls, they increase consumption spending. These users though do not display any of the characteristics that are typically associated with individuals that have precautionary savings motives, such as high income volatility, increasing income paths, or being early on the steep part of the income path. Based of these considerations, we label the mechanism we document in this paper *perceived* precautionary savings motive.

6.3 Alternative Explanations and Channels

We discuss a set of possible interpretations for our results. First, it is implausible the mobile overdraft facility loosens financial constraints on the side of users, because the users with the highest fraction of cash deposits over income react the most to the introduction of the overdraft facility. If anything, the consumption behavior of the most liquidity-constrained users does not change at all.

Second, access to the overdraft facility might help users smooth consumption in case of growing income paths. Our baseline results hold for both young and old users and life-cycle consumption patterns are unlikely to matter for older users.[7] To directly test for this mechanism, we also study the income profiles around the overdraft activation and do not find any evidence of different income growth across bins of deposits to income.

Third, the facility might free up liquid resources users were keeping in their bank accounts due to precautionary motives and potential unexpected future income shocks. In this vein, even models of buffer-stock savings that allow for both impatience and precautionary savings motives might explain at least in part our results. Our heterogeneity results do not seem fully consistent with this interpretation for a set of reasons. Income uncertainty decreases with age and the heterogeneity results by age we discuss above are not consistent with this form of precautionary-savings motive. Moreover, we find the pre-activation volatility of income flows does not predict reaction to the availability of credit. Also, as discussed above, users close to the liquidity constraint do not react, whereas those farther away from the constraint react the most, and the buffer-stock interpretation predicts the opposite pattern.

A fourth potential explanation is a buffer-stock model with durable consumption similar to the one Aydin (2015) studies. In line with this model, we find users mainly spend on durable goods and large-ticket items after activation of the overdraft facility. At the same time, a set of results suggest this interpretation cannot fully explain the results in our setting. In addition to the facts that the least liquidity-constrained users react most and that we do not see any differential reaction based on pre-activation income volatility, we find users that react the most on average do not tap into negative deposits and hence *de facto* never use the facility.

A fifth interpretation we consider is present-biased preferences – the fact individuals discount the distant future by more than they discount the immediate future. This interpretation by itself is unlikely to explain all our results, because if the individuals that react the most to overdraft activation were present-biased, they would have consumed out of their higher deposits even *before* activating the overdraft facility, which we do not observe in the data.

Contrary to the alternative explanations we have discussed in this section, the perceived precautionary savings channel is consistent with the baseline facts we document as well as with the fact that users at the top of the distribution by deposits over income flows react the most

[7]Note that our sample is younger than the average population, yet we do detect an age range of 25 years between the bottom and top quintiles of the distribution by age.

to the activation of the overdraft facility.

7 Conclusion

We study the consumption response to the introduction of an overdraft facility on a FinTech app. The average user increases his consumption spending over income by 4 percentage points at impact relative to similar users that have access to the overdraft facility at a later point in time. The increase in consumption is permanent and we observe a reallocation of consumption from discretionary to non-discretionary expenses. For identification, we exploit a sharp regression discontinuity design eploiting different sizes of overdraft based on income and confirm our baseline findings.

When we study heterogeneity in the response by observables, we observe a similar response for young and old users, for users with low and high income volatility, and for users with high and low future income growth. When we split the sample based on the ratio of deposits over inflows, we find instead a large consumption response for users with high liquid savings, whereas the users with the least liquid savings do not react at all to the provision of the overdraft facility. These results are not fully consistent with myopic consumers, models with financial constraints, buffer stock models (with durables) and present bias and the canonical life-cycle permanent income model.

Additional results, instead, suggest a *perceived* precautionary savings channel is at work for users with high deposits over inflow. Before the facility is available, they perceive high income risk or future expenses and consequently save. Once the overdraft facility is available, which acts like an insurance against future shortfalls, they increase their consumption substantially but barely use the overdraft facility, pay low fees and interests, and their credit ratings are less likely to worsen relative to other users.

Our findings open exciting new avenues for future research. What are the microfoundations of perceived precautionary savings motives? In particular, does this attitude results from biased beliefs about the likelihood of future negative states of the world or is it instead consistent with the neoclassical consumption model in a setting in which consumers have high risk aversion? In terms of policy and real-world applications, do perceived precautionary savings change the effectiveness of conventional fiscal policy such as tax rebates? And could policies be designed to insure perceived precautionary savers in bad times and nudge them to spend their cash in times in which higher aggregate demand is needed?

References

Ausubel, L. M. (1991). The Failure of Competition in the Credit Card Market. *The American Economic Review*, 50–81.

Aydin, D. (2015). The Marginal Propensity to Consume Out of Liquidity: Evidence from Random Assignment of 54,522 Credit Lines. Stanford University Working Paper.

Bain (2017). Evolving the Customer Experience in Banking. https://www.bain.com/insights/evolving-the-customer-experience-in-banking, Accessed on January 3, 2019.

Baron, M. and W. Xiong (2017). Credit Expansion and Neglected Crash Risk. *The Quarterly Journal of Economics 132*(2), 713–764.

Barreca, A. I., M. Guldi, J. M. Lindo, and G. R. Waddell (2011). Saving Babies? Revisiting the Effect of Very Low Birth Weight Classification. *The Quarterly Journal of Economics 126*(4), 2117–2123.

Barreca, A. I., J. M. Lindo, and G. R. Waddell (2016). Heaping-Induced Bias in Regression Discontinuity Designs. *Economic Inquiry 54*(1), 268–293.

Bartlett, R. P., A. Morse, R. H. Stanton, and N. E. Wallace (2018). Consumer-Lending Discrimination in the Era of FinTech. University of California Berkeley, Working Paper.

Calonico, S., M. D. Cattaneo, M. H. Farrell, and R. Titiunik (2018). Regression Discontinuity Designs Using Covariates. The Review of Economics and Statistics (Forthcoming).

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica 82*(6), 2295–2326.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2015). Optimal Data-Driven Regression Discontinuity Plots. *Journal of the American Statistical Association 110*(512), 1753–1769.

Cameron, A. and D. Miller (2015). A Practitioner's Guide to Cluster-Robust Inference. *Journal of Human Resources 50*(2), 317–372.

Campbell, J. Y. and N. G. Mankiw (1989). Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. *NBER Macroeconomics Annual 4*, 185–216.

Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics 112*(1), 1–55.

Cattaneo, M. D., N. Idrobo, and R. Titiunik (2020). *A Practical Introduction to Regression Discontinuity Designs*. Elements in Quantitative and Computational Methods for the Social Sciences. Cambridge University Press.

Cattaneo, M. D., M. Jansson, and X. Ma (2017). Simple Local Polynomial Density Estimators. University of Michigan, Working Paper.

Cheng, M.-Y., J. Fan, and J. S. Marron (1997). On Automatic Boundary Corrections. *The Annals of Statistics 25*(4), 1691–1708.

Claessens, S., J. Frost, G. Turner, and F. Zhu (2018). Fintech Credit Markets around the World: Size, Drivers and Policy Issues. BIS Quarterly Review.

Deaton, A. (1991). Saving and Liquidity Constraints. *Econometrica 59*, 1221–1248.

Di Maggio, M. and A. Kermani (2017). Credit-Induced Boom and Bust. *The Review of Financial Studies 30*(11), 3711–3758.

Dobbie, W., A. Liberman, D. Paravisini, and V. Pathania (2018). Measuring Bias in Consumer Lending. National Bureau of Economic Research, Working Paper.

Friedman, M. (1957). The Permanent Income Hypothesis. In *A Theory of the Consumption Function*, pp. 20–37. Princeton University Press.

Gelman, A. and G. Imbens (2018). Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs. *Journal of Business & Economic Statistics 0*(0), 1–10.

Guerrieri, V. and G. Lorenzoni (2017). Credit Crises, Precautionary Savings, and the Liquidity Trap. *The Quarterly Journal of Economics 132*(3), 1427–1467.

Imbens, G. W. and T. Lemieux (2008). Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics 142*(2), 615–635.

ING (2018). International Survey Mobile Banking 2018. https://www.ezonomics.com/ing_international_surveys/mobile-banking-2018, Accessed on January 3, 2019.

Jappelli, T. and L. Pistaferri (2017). *The economics of consumption: theory and evidence*. Oxford University Press.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity* (1), 77–153.

Laibson, D., A. Repetto, and J. Tobacman (2007). Estimating Discount Functions with Consumption Choices over the Lifecycle. National Bureau of Economic Research, Working Paper.

Lee, D. S. and T. Lemieux (2010, June). Regression Discontinuity Designs in Economics. *Journal of Economic Literature 48*(2), 281–355.

Mian, A. and A. Sufi (2015). House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent it from Happening Again. *University of Chicago Press*.

Mian, A., A. Sufi, and E. Verner (2017). Household debt and business cycles worldwide. *The Quarterly Journal of Economics 132*(4), 1755–1817.

Olafsson, A. and M. Pagel (2018). The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software. *The Review of Financial Studies 31*(11), 4398–4446.

Roberts, M. R. and T. M. Whited (2013). Endogeneity in Empirical Corporate Finance. Volume 2 of *Handbook of the Economics of Finance*, pp. 493–572. Elsevier.

Schularick, M. and A. M. Taylor (2012). Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008. *American Economic Review 102*(2), 1029–61.

Stango, V. and J. Zinman (2009). Exponential Growth Bias and Household Finance. *The Journal of Finance 64*(6), 2807–2849.

Tables and Figures

Table 1: Descriptive Statistics

This table reports descriptive statistics for our user-month panel. The sample consists of 40,979 users that received an overdraft between February 2015 and October 2017 and covers each individual's complete transaction history from February 2015 to September 2018. For each variable, we report the number of observations (N), mean, standard deviation (SD), 10% quantile (P10), 25% quantile (P25), median (P50), 75% quantile (P75), and 90% quantile (P90). We define all variables in Appendix A.

	N	Mean	SD	P10	P25	P50	P75	P90
Age [Years]	595,244	33.906	10.021	23.578	26.746	31.326	38.579	48.912
Time Since Account Opening [Years]	595,244	1.297	0.765	0.356	0.684	1.183	1.840	2.385
Female [0/1=Yes]	595,244	0.204	0.403	0.000	0.000	0.000	0.000	1.000
Urban [0/1=(Population $\geq$ 500k)]	595,244	0.522	0.500	0.000	0.000	1.000	1.000	1.000
Overdraft Available [0/1=Yes]	595,244	0.875	0.330	0.000	1.000	1.000	1.000	1.000
Overdraft Enabled [0/1=Yes]	595,244	0.501	0.500	0.000	0.000	1.000	1.000	1.000
Push Notifications Active [0/1=Yes]	595,244	0.418	0.493	0.000	0.000	0.000	1.000	1.000
Overdraft Amount [Euro]	552,189	1,142.965	912.953	250.000	500.000	750.000	1,500.000	2,500.000
Negative Deposits [0/1=Yes]	520,984	0.322	0.467	0.000	0.000	0.000	1.000	1.000
Overdrawn Amount [Euro]	137,814	705.704	676.392	69.000	212.000	500.000	1,001.000	1,504.000
Overdrawn Amount / Overdraft Amount [%]	137,814	61.865	35.876	8.480	27.244	69.200	100.000	100.000
Fees [Euro]	595,244	2.537	6.733	0.000	0.000	0.000	2.000	6.000
Rating [1-6]	462,396	3.059	1.633	1.000	2.000	3.000	4.000	6.000
Consumption [Euro]	595,244	730.908	983.764	0.000	129.000	453.000	982.000	1,716.000
Inflows [Euro]	595,244	2,072.529	7,747.827	50.000	301.000	1,001.000	2,367.000	4,324.000
Cash Withdrawals [Euro]	595,244	239.208	463.382	0.000	0.000	91.000	306.000	640.000
Card Consumption [Euro]	595,244	481.574	748.682	0.000	58.000	259.000	614.000	1173.000
Discretionary [Euro]	595,244	203.245	437.227	0.000	0.000	64.000	230.000	534.000
Non-Discretionary [Euro]	595,244	278.329	479.039	0.000	20.000	143.000	352.000	670.000
Entertainment [Euro]	595,244	14.921	69.867	0.000	0.000	0.000	9.000	35.000
Gastronomy [Euro]	595,244	40.756	110.555	0.000	0.000	0.000	36.000	118.000
Groceries [Euro]	595,244	79.900	118.686	0.000	0.000	33.000	120.000	222.000
Shopping [Euro]	595,244	61.776	188.833	0.000	0.000	0.000	52.000	177.000
Travel [Euro]	595,244	85.792	299.022	0.000	0.000	0.000	60.000	220.000
Consumption / Inflows [%]	521,352	54.355	53.726	0.250	15.385	40.069	76.787	123.500
Card Consumption / Inflows [%]	521,369	34.540	38.274	0.000	6.507	22.383	48.696	87.096
Cash Withdrawals / Inflows [%]	521,537	14.453	20.861	0.000	0.000	5.834	20.579	42.839
Discretionary / Non-Discretionary [%]	455,694	80.571	103.907	0.000	9.890	42.058	107.286	218.537
Entertainment / Card Consumption [%]	487,775	2.218	5.025	0.000	0.000	0.000	1.786	7.491
Gastronomy / Card Consumption [%]	487,762	5.928	9.292	0.000	0.000	0.000	8.879	20.247
Shopping / Card Consumption [%]	487,762	7.905	12.849	0.000	0.000	0.000	11.738	28.261
Travel / Card Consumption [%]	487,762	9.572	14.903	0.000	0.000	0.000	14.213	32.915
Big Ticket Expense (>1000 Euro) [0/1=Yes]	595,244	0.431	0.495	0.000	0.000	0.000	1.000	1.000
Big Ticket Expense (>2000 Euro) [0/1=Yes]	595,244	0.220	0.415	0.000	0.000	0.000	0.000	1.000

Table 2: Effect of Mobile Overdrafts on Users' Consumption Behavior

This table provides coefficient estimates of OLS regressions estimating the effect of mobile overdraft facilities on users' consumption behavior (equation (1)). *Consumption* is the sum of users' *Card Consumption* and *Cash Withdrawals* in the given month. *Card Consumption* is the user's total amount of electronic card consumption. *Cash Withdrawals* is the user's total amount of cash withrawals from ATMs in the given month. *Inflows* is the total amount of all incoming transactions a user receives in the given month. *Overdraft Available* is a binary indicator that equals 1 if the user has access to a mobile overdraft in the given month. We report t-statistics based on standard errors double-clustered at the NUTS2 and year-month level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable (×100):	$\frac{\text{Consumption}_t}{\text{Inflows}_{t-1}}$	$\frac{\text{Card Consumption}_t}{\text{Inflows}_{t-1}}$	$\frac{\text{Cash Withdrawals}_t}{\text{Inflows}_{t-1}}$
	(1)	(2)	(3)
Overdraft Available$_t$	4.413***	2.952***	1.051***
	(11.30)	(8.83)	(6.29)
Fixed Effects:			
User	Yes	Yes	Yes
NUTS3 × Year-Month	Yes	Yes	Yes
Standard Error Clusters:			
NUTS2	48	48	48
Year-Month	43	43	43
Adjusted R^2	0.254	0.284	0.330
User-Year-Month Observations	517,114	517,149	517,314

Figure 2: Consumption Pattern around Mobile Overdraft Availability
This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the effect of mobile overdrafts on users' consumption behavior. We estimate model (1) from Table 2 but replace the *Overdraft Available* indicator with separate time dummies, each marking a one-month period (except for event period t-1).

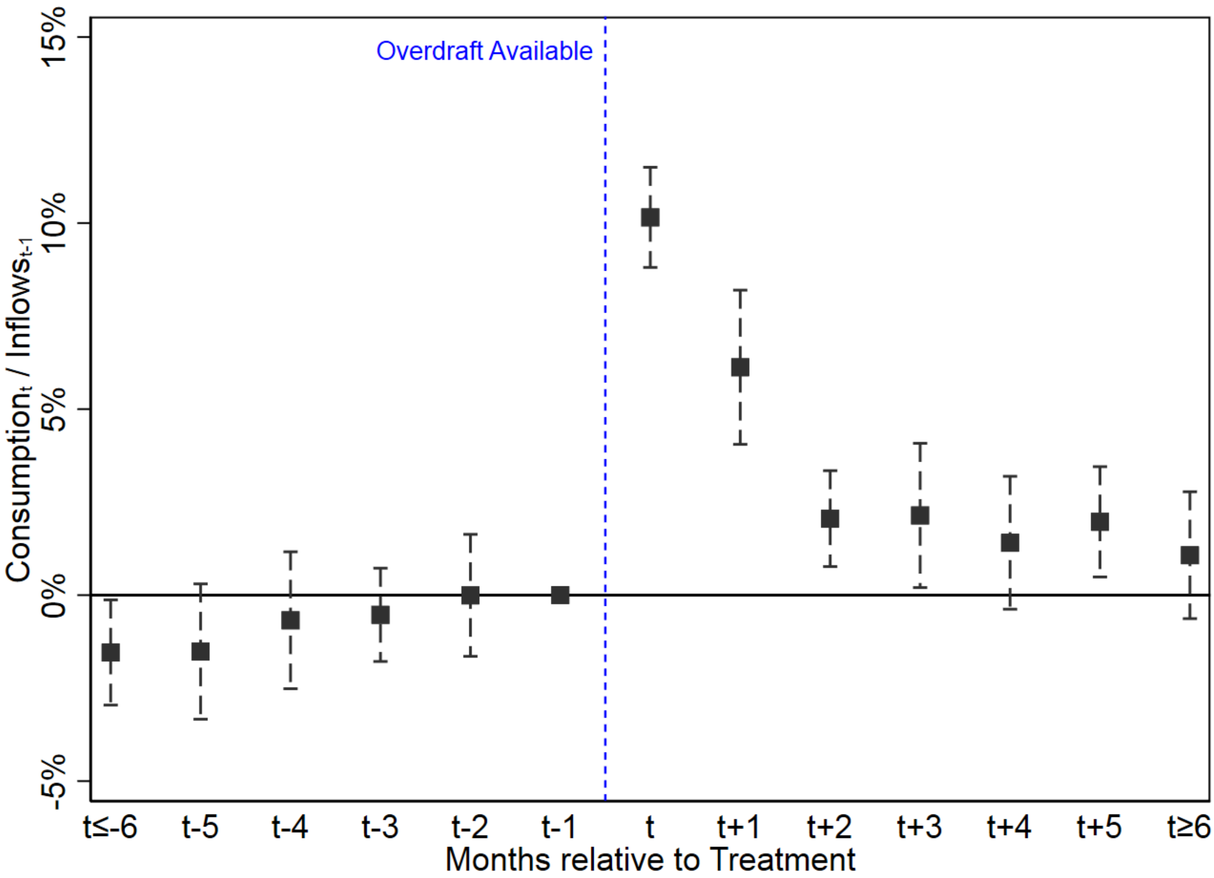

Figure 3: Patterns of Card Consumption and Cash Withdrawals around Mobile Overdraft Availability This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the effect of mobile overdrafts on users' card consumption and cash withdrawals. We estimate models (2) and (3) from Table 2 but replace the *Overdraft Available* indicator with separate time dummies, each marking a one-month period (except for event period t-1).

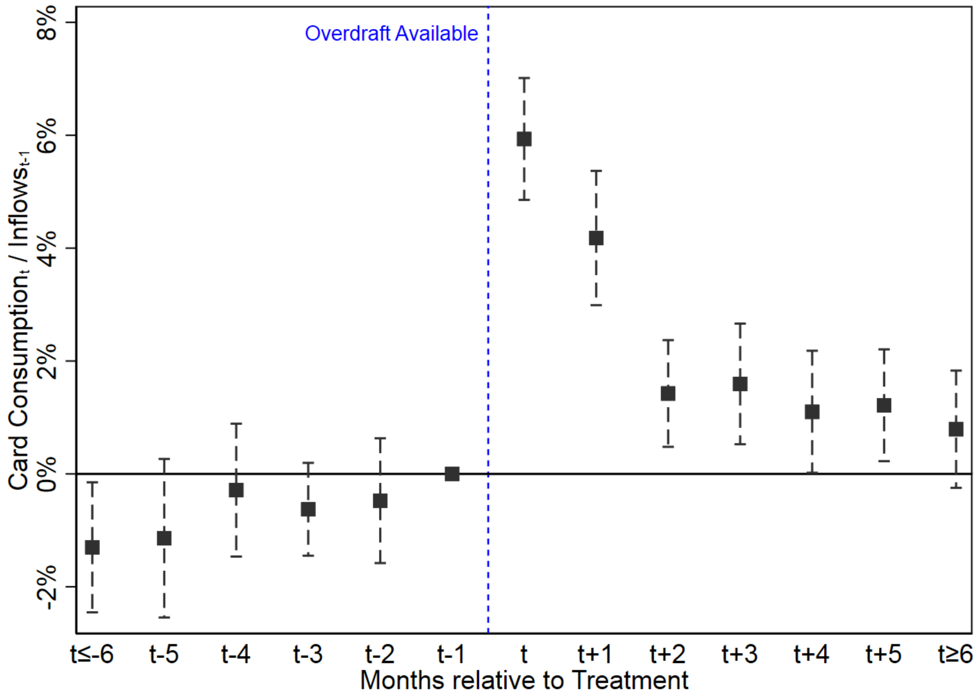

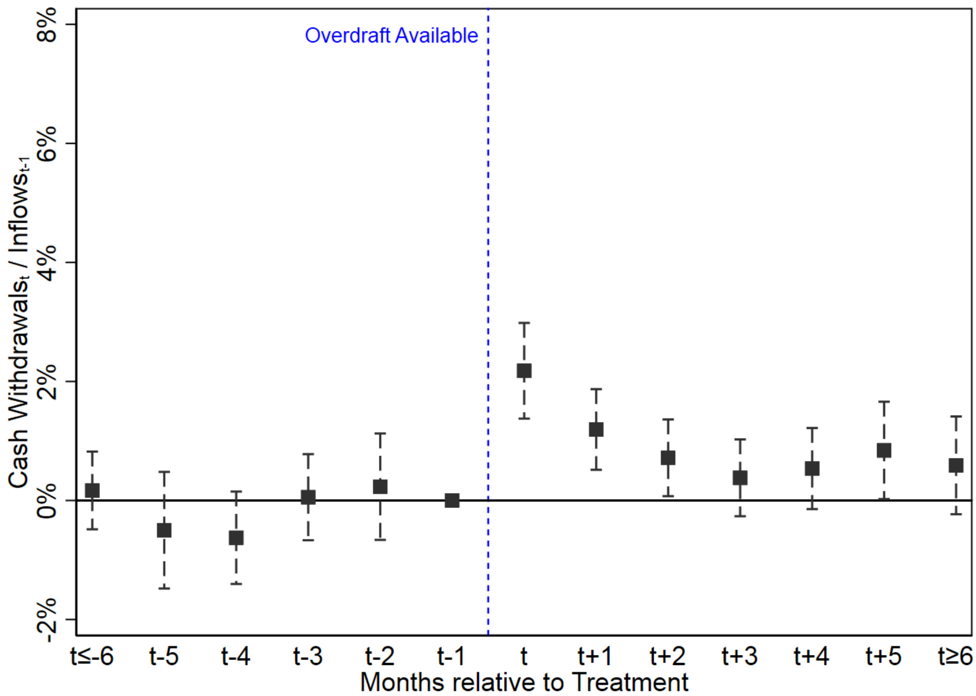

Table 3: Effect of Mobile Overdrafts on Users' Credit Risk

This table reports coefficient estimates of OLS regressions estimating the effect of mobile overdrafts on users' credit risk (equation (2)). *Rating* is a discrete variable that ranges from 1 (highest rating) to 6 (lowest rating). *Overdraft Available* is a binary indicator that equals 1 if the user has access to a mobile overdraft in the given month. The sample in column (1) consists of all users that successfully applied for an overdraft. In column (2), we only focus on users that experienced at least 1 rating change over our sample period. We report t-statistics based on standard errors double-clustered at the NUTS2 and year-month level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable:	Rating_t	
	(1)	(2)
	All Users	Users with Rating Changes
$\text{Overdraft Available}_t$	0.026**	0.992***
	(2.81)	(3.21)
Fixed Effects:		
User	Yes	Yes
NUTS3 × Year-Month	Yes	Yes
Standard Error Clusters:		
NUTS2	48	
Year-Month	24	22
Adjusted R^2	0.998	0.698
User-Year-Month Observations	259,705	622

Figure 4: Credit Risk Pattern around Mobile Overdraft Availability
This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the effect of mobile overdrafts on users' credit risk. We estimate model (1) from Table 3 but replace the *Overdraft Available* indicator with separate time dummies, each marking a one-month period (except for event period t).

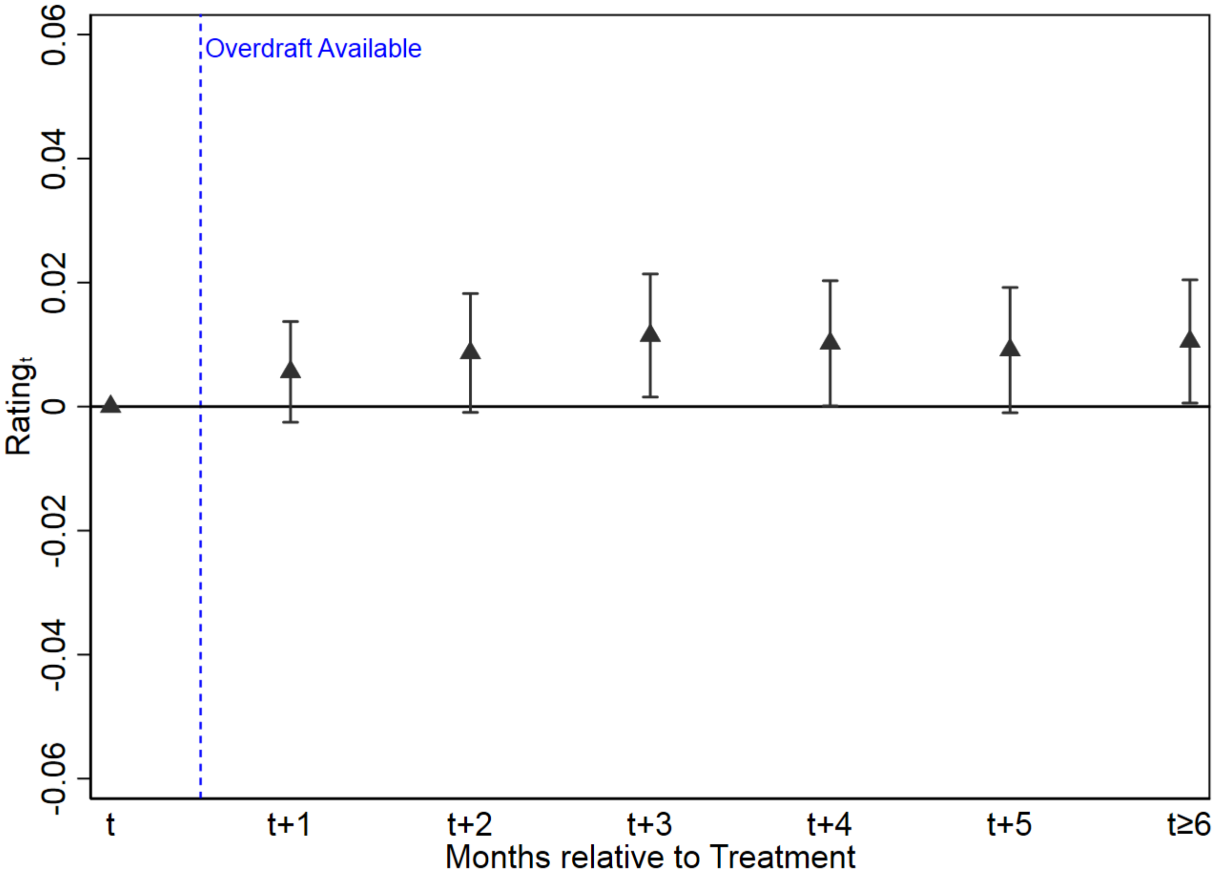

Table 4: Mobile Overdraft Availability on Mobile Overdraft Usage

This table reports coefficient estimates of OLS regressions estimating the association between mobile overdraft availability and mobile overdraft usage (equation (3)). *Overdraft Enabled* is binary indicator that equals 1 if the user enabled the overdraft and *Negative Deposits* is a dummy variable equal to 1 if the individual actively uses the overdraft and has a negative account balance. *Overdraft Available* is a binary indicator that equals 1 if the user has access to a mobile overdraft in the given month. *Overdraft Amount* is the users' maximum overdraft limit. We report t-statistics based on standard errors double-clustered at the NUTS2 and year-month level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable:	Overdraft Enabled		Negative Deposits	
	(1) Extensive Margin	(2) Intensive Margin	(3) Extensive Margin	(4) Intensive Margin
Overdraft Available$_t$	0.708***		0.482***	
	(28.07)		(29.26)	
Log(Overdraft Amount$_t$)		0.052***		0.059***
		(10.17)		(16.10)
Fixed Effects:				
User	Yes	Yes	Yes	Yes
NUTS3 × Year-Month	Yes	Yes	Yes	Yes
Standard Error Clusters:				
NUTS2	48	48	48	48
Year-Month	35	35	35	35
Adjusted R^2	0.929	0.933	0.551	0.551
User-Year-Month Observations	532,909	518,108	532,909	518,108

Table 5: Effect of Mobile Overdrafts on Fees

This table reports coefficient estimates of OLS regressions estimating the effect of mobile overdrafts on users' quarterly account fees (equation (4)). *Fees* is the total amount of fees that the bank deducts from the individual's checking account in quarter q, normalized by the user's contemporaneous account inflows. *Overdraft Available* is a binary indicator that equals 1 if the user had access to a mobile overdraft in at least one month of the given quarter. *Overdraft Enabled* is binary indicator that equals 1 if the user enabled the overdraft in at least one month of the given quarter. *Overdraft Amount* is the users' average maximum overdraft limit during the quarter. We report t-statistics based on standard errors clustered at the NUTS2 level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable (×100):	$\frac{\text{Fees}_q}{\text{Inflows}_q}$		
	(1)	(2)	(3)
Overdraft Available$_{q-1}$	0.353***		
	(4.07)		
Overdraft Enabled$_{q-1}$		0.493**	
		(2.25)	
Log (Overdraft Amount$_{q-1}$)			0.263**
			(2.44)
Fixed Effects:			
User	Yes	Yes	Yes
NUTS3 × Year-Quarter	Yes	Yes	Yes
Standard Error Clusters:			
NUTS2	48	48	48
Adjusted R^2	0.088	0.088	0.120
User-Year-Quarter Observations	219,832	219,832	174,396

Figure 5: Fee Pattern around Mobile Overdraft Availability
This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the effect of mobile overdrafts on users' quarterly account fees. We estimate model (1) from Table 5 but replace the *Overdraft Available* indicator with separate time dummies, each marking a one-quarter period (except for event period q).

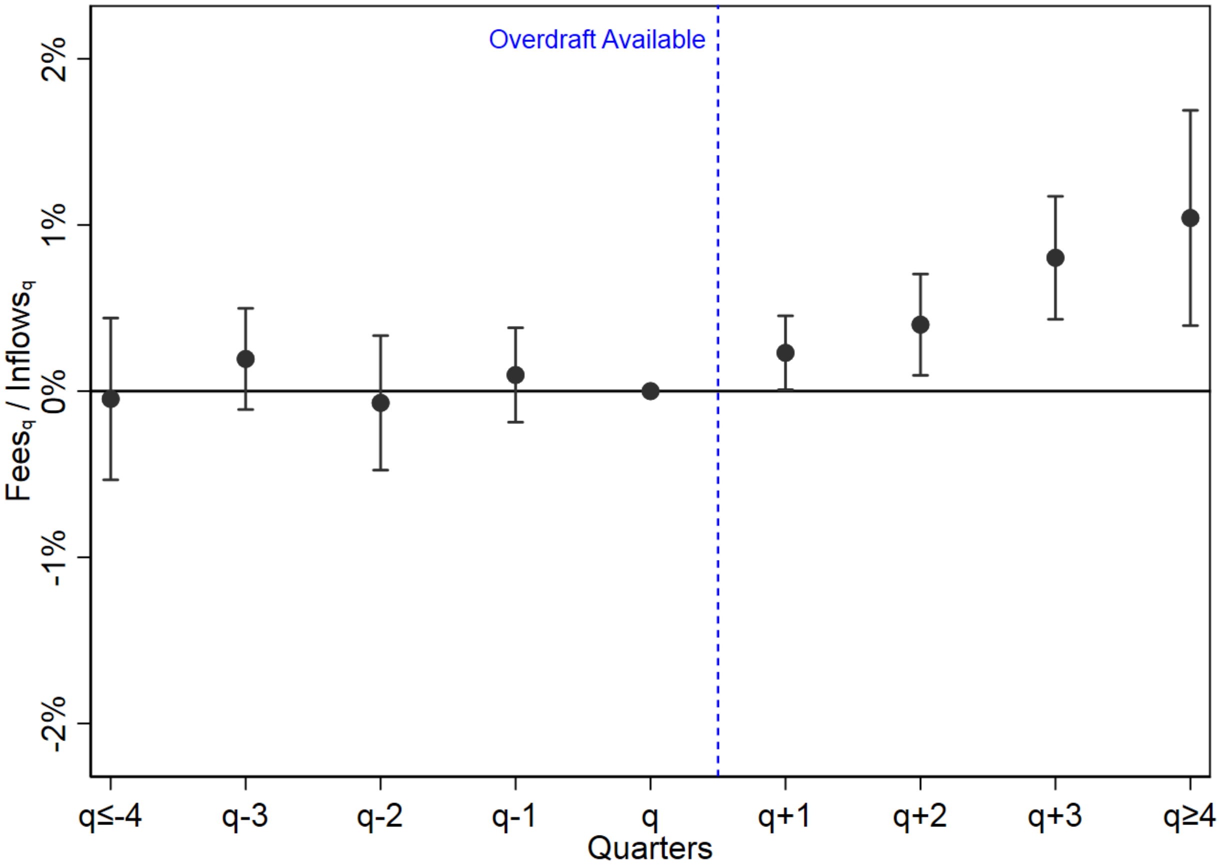

Figure 6: Treatment Assignment in Sharp RD Analysis
This figure illustrates how we assign users to treatment and control group in our regression discontinuity analysis based on discrete rounding thresholds embedded in the bank's credit risk model. Panel A visualizes the rounding logic of the overdraft allocation algorithm. Panel B plots users' treatment probability for different values of our forcing variable X_i.

Panel A: Rounding Logic of Overdraft Allocation Algorithm

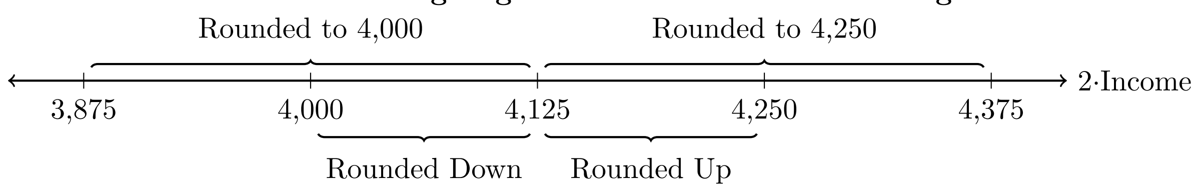

Panel B: Visualization of Sharp Treatment

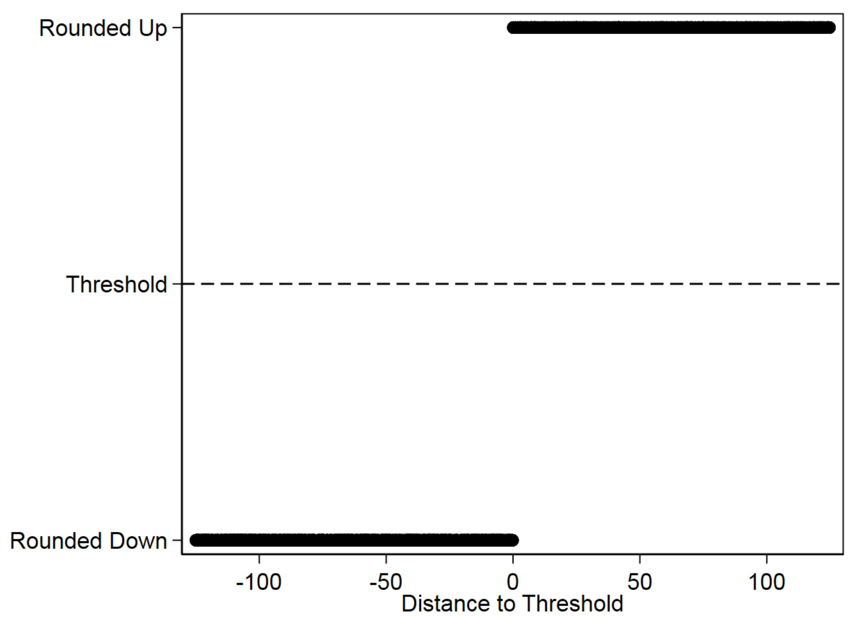

Figure 7: Distribution of Forcing Variable around Rounding Threshold
This figure provides graphical evidence for our treatment manipulation tests in Section **??**. Panel A plots the number of users and Panel B reports the local polynomial density estimate by Cattaneo et al. (2017) for different values of our running variable X_i around the rounding threshold.

Panel A: Histogram of Forcing Variable

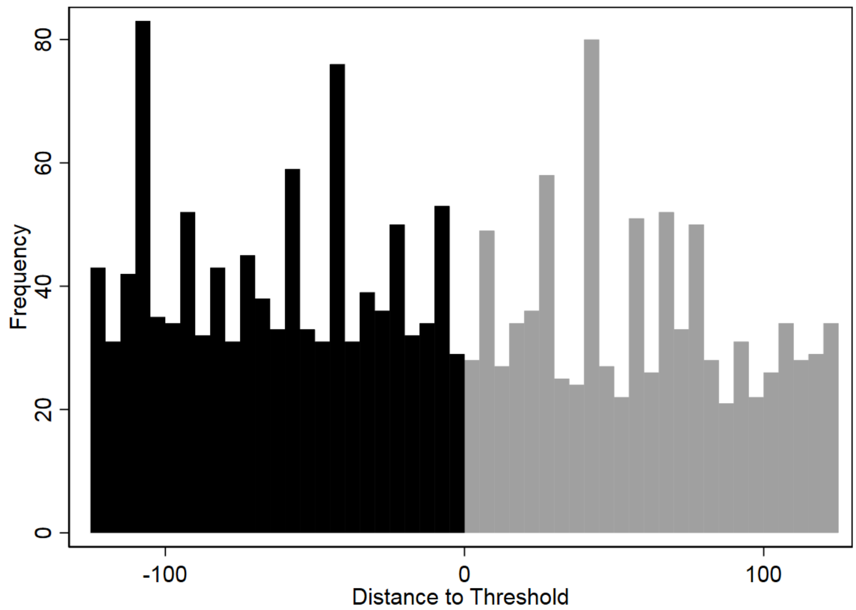

Panel B: Local Polynomial Density Estimate

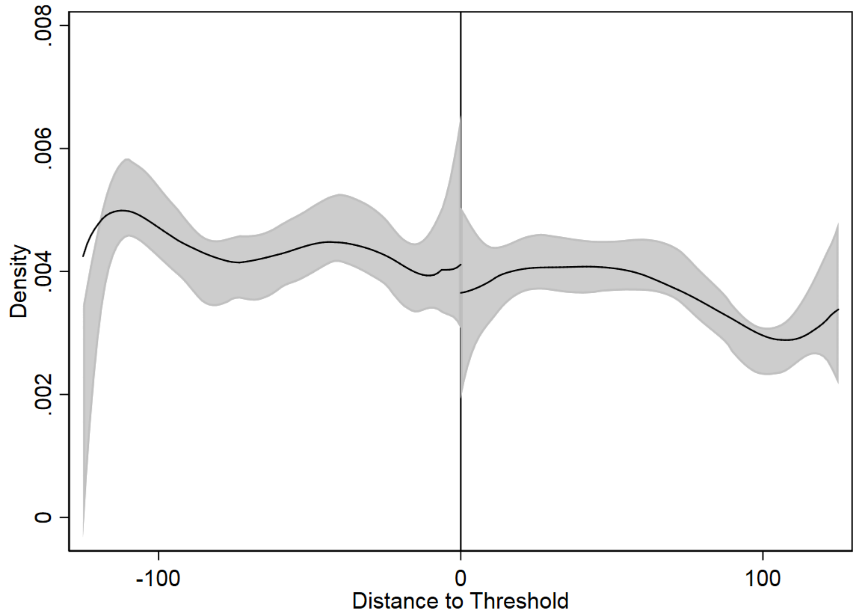

Table 6: Treatment Manipulation Tests

This table reports the results of treatment manipulation tests using the local polynomial density estimator by Cattaneo et al. (2017). $T_q(h_p)$ denotes the q-th order local polynomial test with bandwidth h_p. "Bandwidth" is the mean-squared-error (MSE) optimal bandwidth, "Effective N" is the effective sample size on each side of the threshold, and "T" is the two-sided test statistic with corresponding p-value. The tests in the first three rows allow for different bandwidths while the tests in the last three rows impose a common bandwidth on both sides of the threshold.

		Bandwidth		Effective N		Test	
		Left	Right	Left	Right	T	p-value
$h_- \neq h_+$	$T_2(\hat{h}_1)$	35.05	47.84	273	376	-1.10	0.27
	$T_3(\hat{h}_2)$	57.00	55.26	456	412	-1.19	0.23
	$T_4(\hat{h}_3)$	53.48	59.63	432	460	-0.15	0.88
$h_- = h_+$	$T_2(\hat{h}_1)$	34.51	34.51	269	254	-0.52	0.60
	$T_3(\hat{h}_2)$	125.77	125.77	1045	875	0.56	0.58
	$T_4(\hat{h}_3)$	69.45	69.45	569	533	-0.91	0.36

Table 7: Descriptive Statistics for Users around the Rounding Threshold

This table provides descriptive user characteristics for individuals above ("Rounded Up") and below ("Rounded Down") the RD rounding threshold. For each variable, we report the mean, standard deviation (SD), and median (P50). In the last two columns, we test for differences in means across both types of users. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

	Rounded Up			Rounded Down			Difference in Means	
	Mean	SD	P50	Mean	SD	P50	Diff. Mean	t-Stat.
Age [Years]	32.797	8.725	30.998	33.179	9.444	30.913	0.382	(0.92)
Female [0/1=Yes]	0.184	0.388	0.000	0.170	0.376	0.000	-0.014	(-0.78)
Time Since Account Opening [Years]	0.983	0.396	0.925	0.977	0.395	0.931	-0.006	(-0.31)
Rating [1-6]	2.429	1.403	2.000	2.389	1.278	2.000	-0.040	(-0.65)
Inflows$_{t-3:t-1}$ [Euro]	2,380.941	1,756.363	1,948.333	2,414.163	3,861.670	1,804.000	33.222	(0.25)
Consumption$_{t-3:t-1}$ [Euro]	978.117	777.105	826.000	924.434	755.833	762.000	-53.683	(-1.53)
User Observations	875			1,045			1,920	

Figure 8: Regression Discontinuity Plots

This figure provides graphical evidence for the sharp discontinuity in users' overdraft limits and consumption growth rates at the rounding threshold. In Panel A, we plot the income-normalized overdraft limit that users receive at the treatment date. In Panel B, we plot users' consumption growth rate, which we define as the difference in average consumption three months before and after the treatment, normalized by the average account inflows three months prior to the overdraft application. We aggregate our data into 16 disjoint bins, calculate the average value, plot this value above the midpoint of each bin, and separately fit 2 linear regressions through all observations on each side of the rounding threshold.

Panel A: Overdraft Amounts

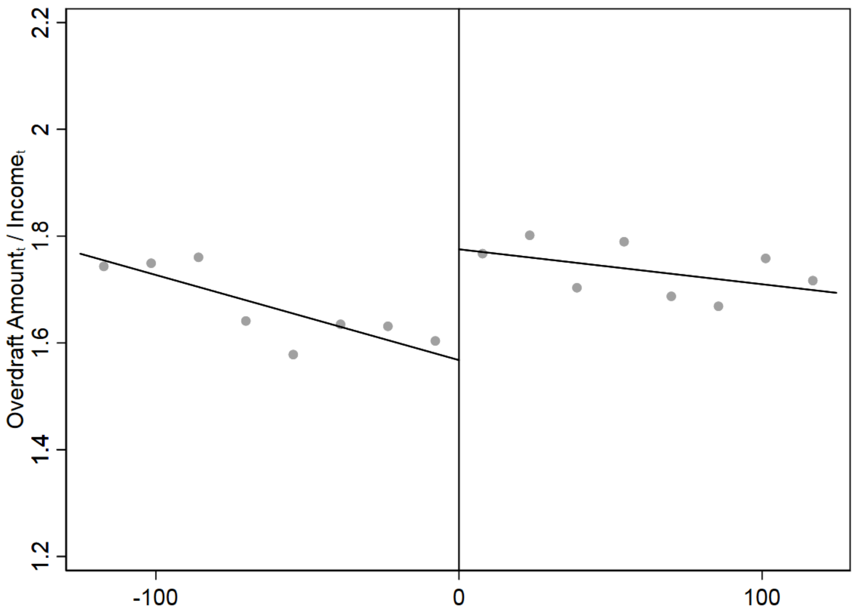

Panel B: Consumption Growth

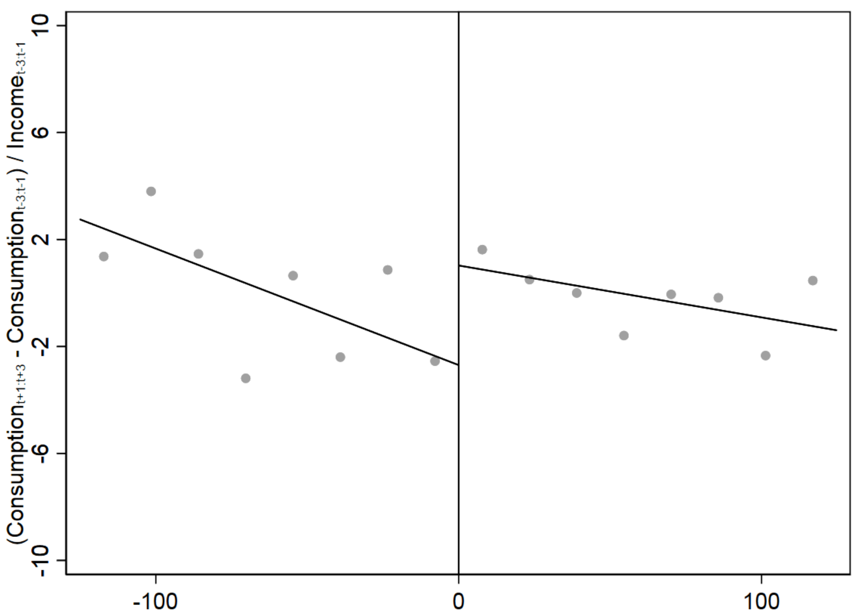

Table 8: Consumption Growth around Rounding Threshold

This table presents non-parametric estimates for the RD treatment effect of a 250 Euro higher overdraft amount on users' consumption behavior. The dependent variable is the difference in average consumption three months before and after the treatment, normalized by the average account inflows three months prior to the overdraft application. We residualize users' consumption growth rate with country × year-month fixed effects to ensure that we compare treated and control users from the same European country at a similar point in time. We only use polynomials of order 1 and 2 to avoid overfitting issues (Gelman and Imbens (2018)), apply weights from a triangular kernel because it is the mean squared error (MSE) minimizing choice (Cheng et al., 1997), and employ the MSE-optimal bandwidth selection procedure recommended by Calonico et al. (2014). We report both conventional and robust RD estimates (Calonico et al., 2014, 2018). In columns (1) and (2), we do not add any covariates. In columns (3) and (4), we control for *Age*, gender (*Female*), and *Time Since Account Opening*. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable (×100):	$\frac{\text{Consumption}_{t+1:t+3} - \text{Consumption}_{t-3:t-1}}{\text{Inflows}_{t-3:t-1}}$			
	(1)	(2)	(3)	(4)
Conventional	10.93**	12.89**	10.87**	12.59**
	(2.36)	(2.45)	(2.38)	(2.49)
Robust	12.82**	14.55**	12.80**	14.27**
	(2.44)	(2.52)	(2.47)	(2.57)
Covariates	No	No	Yes	Yes
User Observations	1,906	1,906	1,906	1,906
Order Local Polynomial (p)	1	2	1	2
Order Bias (q)	2	3	2	3
Bandwidth Left	23.90	40.01	23.14	40.68
Bandwidth Right	23.90	40.01	23.14	40.68
Effective Observations Left	163	301	161	308
Effective Observations Right	165	284	159	289

Figure 9: RD Consumption Effect over Time

This figure visualizes the RD consumption effect over time. We estimate model (1) from Table 8, using 4 separate consumption growth indicators as outcome variable, each defined as the individual's average consumption in a 3-month period following the RD treatment minus the average consumption in the 3 months before the overdraft application, normalized by the user's average account inflows in the quarter prior to the treatment. We plot RD point estimates and their 95% confidence intervals.

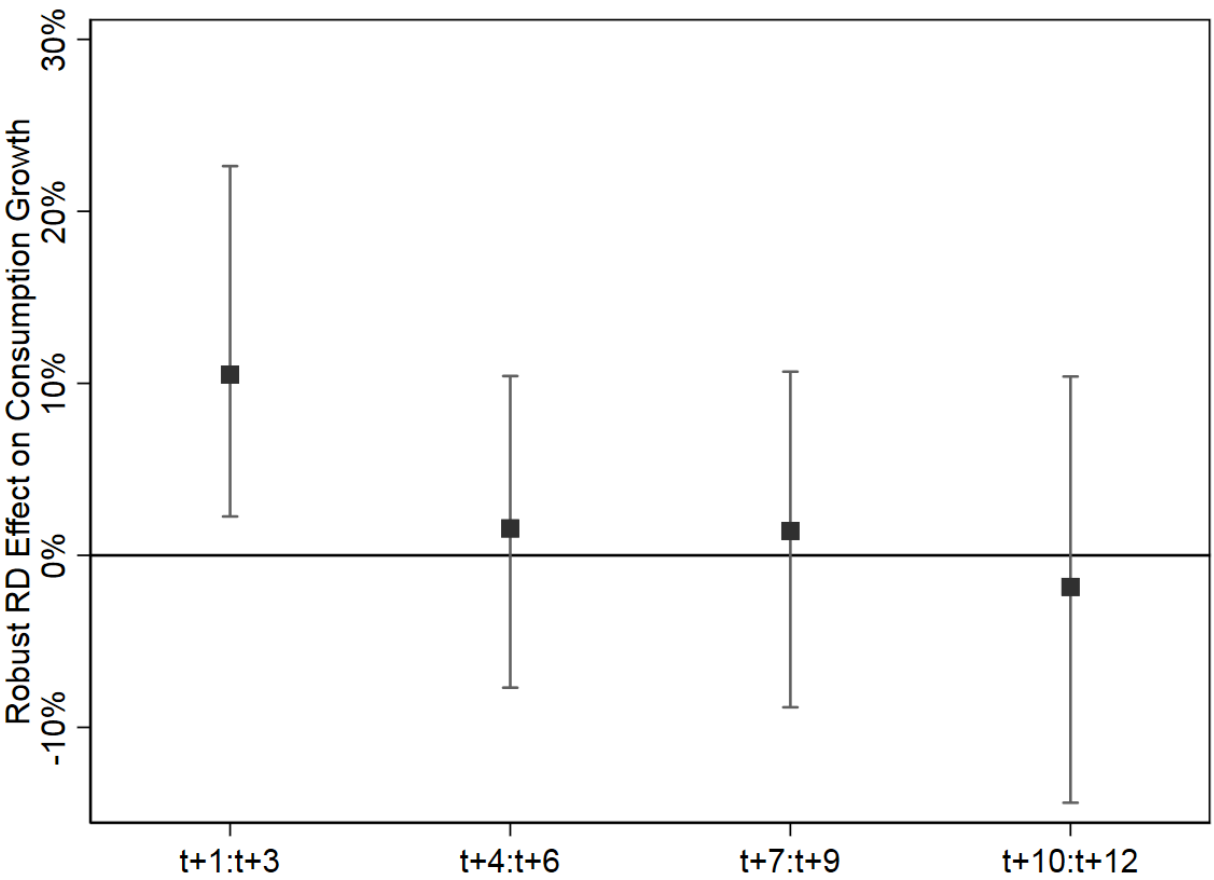

Table 9: Consumption Reallocation Effects of Mobile Overdraft Facilities
This table provides coefficient estimates of OLS regressions estimating the effect of mobile overdrafts on the composition of users' electronic consumption expenditures. *Discretionary* is the sum of users' monthly spending on *Entertainment*, *Shopping*, *Gastronomy*, and *Travel*. *Non-Discretionary* consumption equals *Card Consumption* minus *Discretionary* spending. *Big Ticket Expense (x)* is an indicator variable equal to one if the user purchased at least one item with a transaction amount exceeding x Euros in the given month. *Overdraft Available* is a binary indicator that equals 1 if the user has access to a mobile overdraft in the given month. We report t-statistics based on standard errors double-clustered at the NUTS2 and year-month level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Panel A: Consumption Categories

Dependent Variable (×100):	$\frac{\text{Discretionary}_t}{\text{Non-Discretionary}_t}$	$\frac{\text{Entertainment}_t}{\text{Card Consumption}_t}$	$\frac{\text{Shopping}_t}{\text{Card Consumption}_t}$	$\frac{\text{Gastronomy}_t}{\text{Card Consumption}_t}$	$\frac{\text{Travel}_t}{\text{Card Consumption}_t}$
	(1)	(2)	(3)	(4)	(5)
Overdraft Available$_t$	2.025***	0.073**	0.302***	0.096*	0.649***
	(3.75)	(2.12)	(3.22)	(1.91)	(5.63)
Fixed Effects:					
User	Yes	Yes	Yes	Yes	Yes
NUTS3 × Year-Month	Yes	Yes	Yes	Yes	Yes
Standard Error Clusters:					
NUTS2	48	48	48	48	48
Year-Month	44	44	44	44	44
Adjusted R^2	0.164	0.305	0.172	0.296	0.168
User-Year-Month Observations	451,149	483,640	483,578	483,596	483,542

Panel B: Probability of Big Ticket Expenses

Dependent Variable:	Big Ticket Expense (1000)	Big Ticket Expense (2000)
	(1)	(2)
Overdraft Available$_t$	0.019***	0.012***
	(5.10)	(4.83)
Fixed Effects:		
User	Yes	Yes
NUTS3 × Year-Month	Yes	Yes
Standard Error Clusters:		
NUTS2	48	48
Year-Month	44	44
Adjusted R^2	0.578	0.568
User-Year-Month Observations	591,997	591,997

Figure 10: Consumption Reallocation around Mobile Overdraft Availability
This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the effect of mobile overdrafts on the composition of users' electronic consumption expenditures. We estimate model (1) from Panel A in Table 9 but replace the *Overdraft Available* indicator with separate time dummies, each marking a one-month period (except for event period t-1).

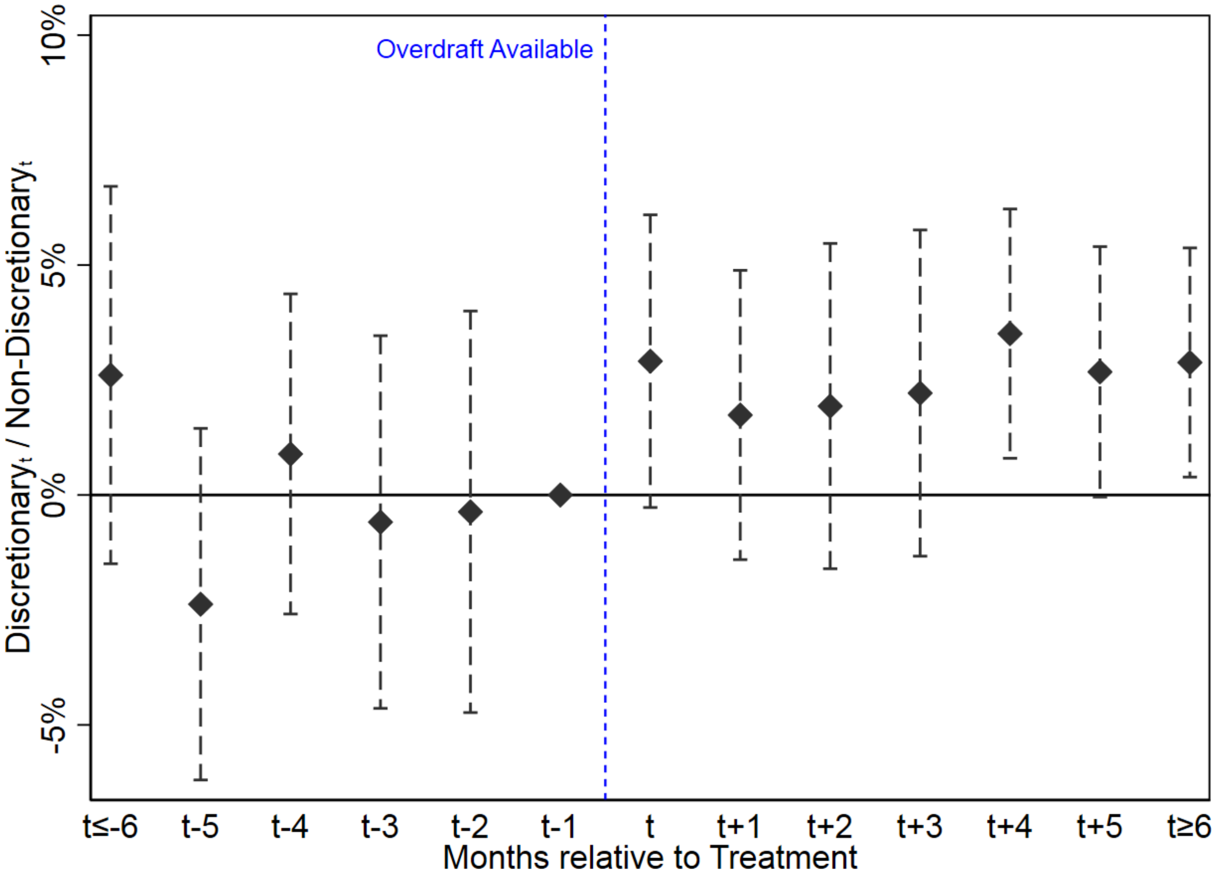

Figure 11: Discretionary Spending around Mobile Overdraft Availability
This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the effect of mobile overdrafts on the composition of users' discretionary consumption behavior. We estimate models (2)-(5) from Panel A in Table 9 but replace the *Overdraft Available* indicator with separate time dummies, each marking a one-month period (except for event period t-1).

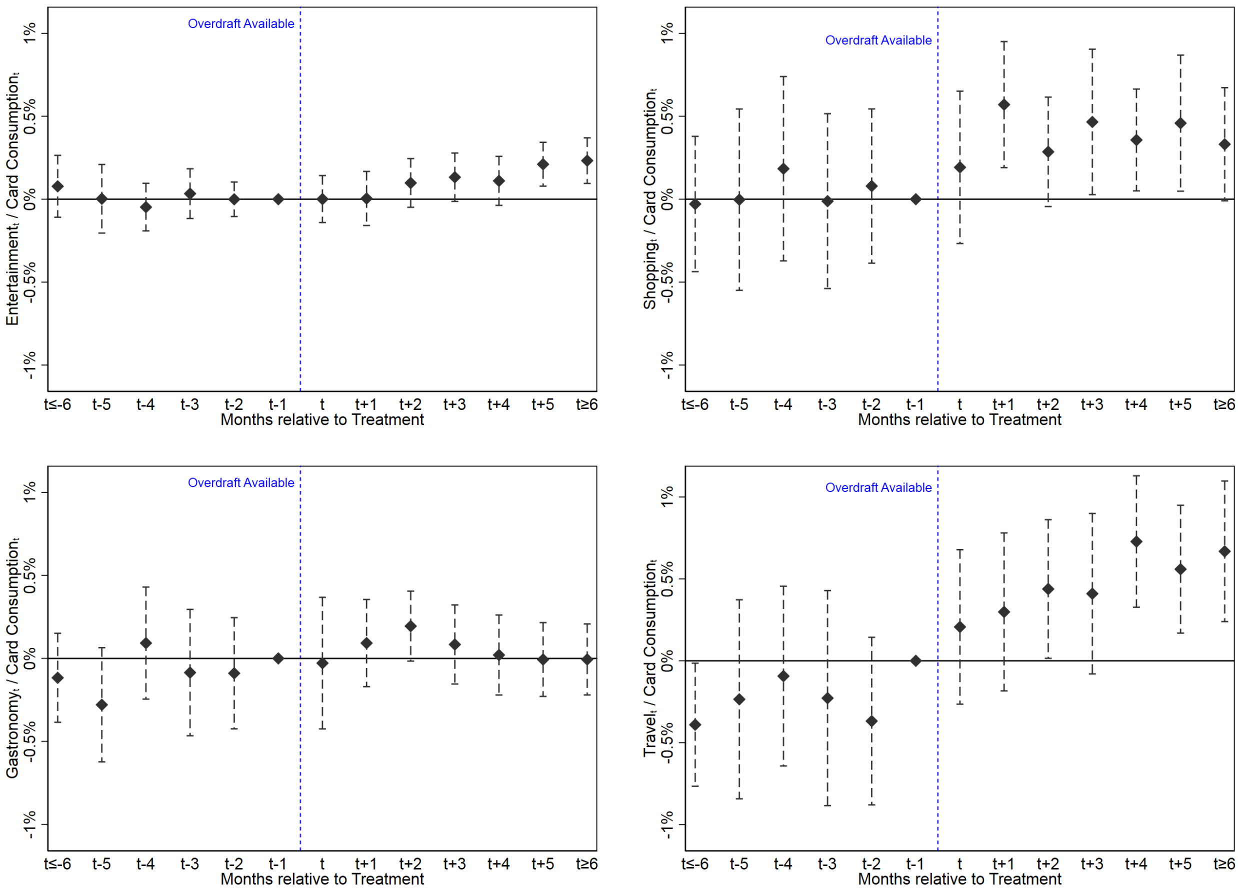

Figure 12: Big Ticket Purchases around Mobile Overdraft Availability
This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the effect of mobile overdrafts on the probability of users making big ticket purchases. We estimate models (1) and (2) from Panel B in Table 9 but replace the *Overdraft Available* indicator with separate time dummies, each marking a one-month period (except for event period t-1).

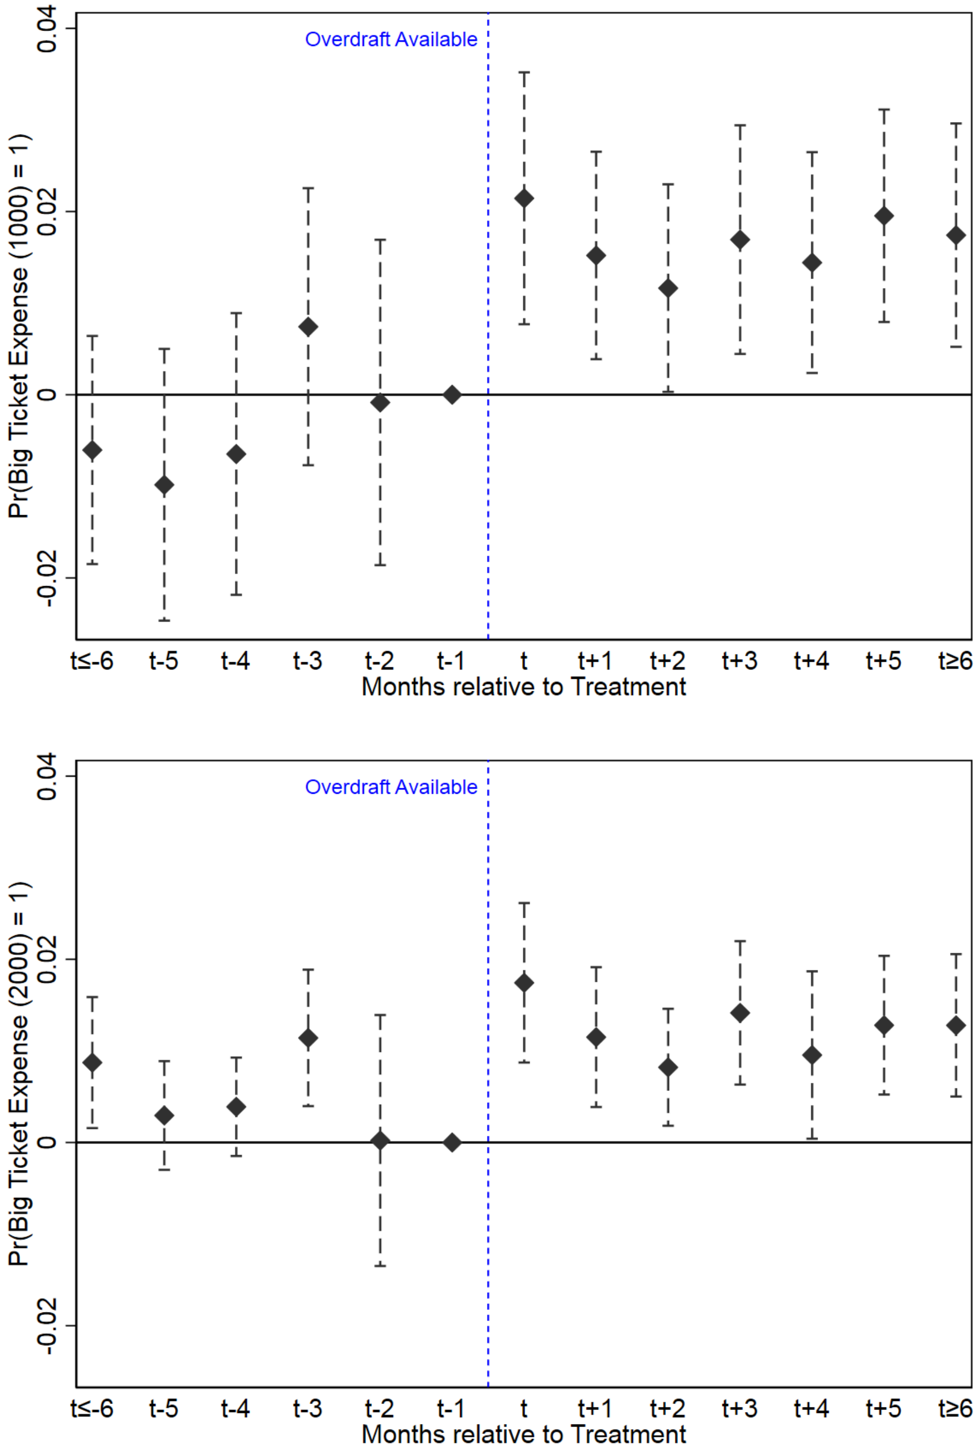

Table 10: Cross-Sectional Heterogeneity in Consumption Response

This table provides coefficient estimates of OLS regressions examining cross-sectional heterogeneity in users' consumption response to the mobile overdraft. We estimate model (1) from Table 3 but add interaction terms of *Overdraft Available* with several cross-sectional user characteristics. *Consumption* is the sum of users' *Card Consumption* and *Cash Withdrawals* in the given month. *Inflows* is the total amount of all incoming transactions a user receives in the given month. *Overdraft Available* is a binary indicator that equals 1 if the user has access to a mobile overdraft in the given month. For each cross-sectional test, we partition individuals into two non-overlapping groups based on the sample median of the corresponding user characteristic at the time users obtain access to the overdraft. *Inflows Growth* is the user's average growth of inflows 6 months before to 6 months after the overdraft activation. *Age* is the user's age at the treatment date in years. *Deposits / Inflows* is the ratio of deposits to inflows in the month prior to the overdraft activation. *Time Since Account Opening* indicates how long the user has had a checking account at the bank in years. We report t-statistics based on standard errors double-clustered at the NUTS2 and year-month level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable (×100):	$\frac{\text{Consumption}_t}{\text{Inflows}_{t-1}}$		
	(1)	(2)	(3)
Overdraft Available$_t$	3.762***	3.437***	-0.096
	(6.64)	(11.25)	(-0.13)
Overdraft Available$_t$ * Inflows Growth > Median	-1.917**		
	(-2.33)		
Overdraft Available$_t$ * Age > Median		0.231	
		(0.77)	
Overdraft Available$_t$ * Deposits / Inflows > Median			7.153***
			(7.58)
Fixed Effects:			
User	Yes	Yes	Yes
NUTS3 × Year-Month	Yes	Yes	Yes
Standard Error Clusters:			
NUTS2	48	48	48
Year-Month	43	43	43
Adjusted R^2	0.244	0.247	0.246
User-Year-Month Observations	77,925	265,951	219,442

Figure 13: Cross-Sectional Quintile Plots
This figure illustrates the cross-sectional heterogeneity in users' consumption response to the mobile overdraft. To generate these plots, we take the cross-section of users at their treatment date and assign them into non-overlapping quintiles from lowest (1st quintile) to highest (5th quintile) based on the underlying user characteristic. We then interact each of the 5 quintile indicators with a dummy variable that equals 1 if the user has access to a mobile overdraft in the given month. Vertical bands represent 95% confidence intervals for the point estimates of each quintile. We double cluster standard errors at the NUTS2 and year-month level.

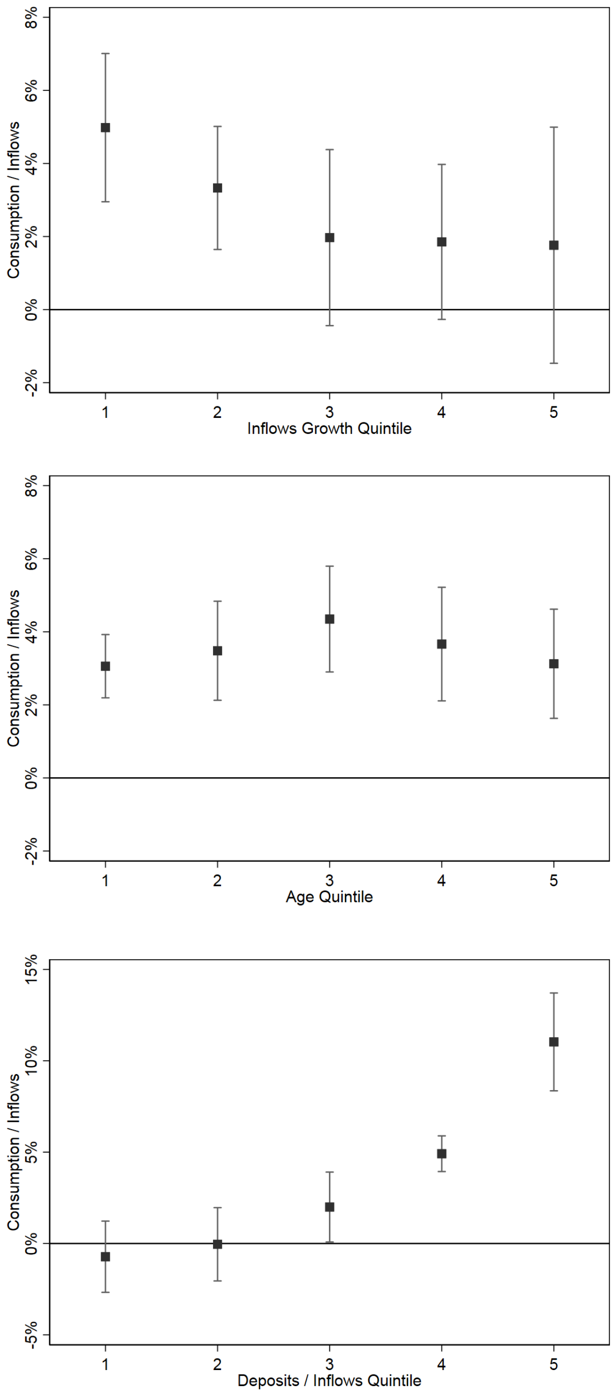

Table 11: User Characteristics for Deposit over Inflows Quintiles

This table provides descriptive statistics for users at the time of overdraft activation, sorted into quintiles based on the ratio deposits over inflows in the month before activation. Subscript t refers to the activation date, while T denotes the last month of each user in our sample. The last two columns test for differences in means across the two groups of users. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

	Deposits / Inflows Quintiles					Difference in Means	
	(1)	(2)	(3)	(4)	(5)	(5)-(1)	t-Stat.
Age_t	31.455	32.212	32.695	33.033	34.910	3.454***	(10.99)
Female	0.194	0.197	0.209	0.240	0.232	0.038***	(3.00)
Time Since Account Opening_t	0.805	0.701	0.705	0.613	0.747	-0.058***	(-4.08)
Deposits_{t-1}	32.596	439.433	1078.545	1909.177	4124.741	4092.145***	(10.65)
$\text{Deposits}_{t-12:t-1}$	316.625	541.457	891.213	1258.738	2362.494	2045.869***	(9.92)
Inflows_{t-1}	1309.255	1815.831	1865.397	2009.052	1321.968	12.713	(0.10)
$\text{Inflows}_{t-12:t-1}$	1188.622	1335.739	1522.682	1383.086	1143.447	-45.175	(-0.41)
$\text{SD}(\text{Inflows}_{t-12:t-1})$	1026.661	1070.432	1129.470	1364.052	1060.108	33.447	(0.26)
Inflows Growth $(\pm\ 3\ \text{Months})_t$	2.497	1.854	2.390	3.555	4.579	2.082	(1.62)
Inflows Growth $(\pm\ 6\ \text{Months})_t$	2.559	1.601	1.645	2.487	4.079	1.520	(1.33)
Positive Consumption Change $(\pm\ 3\ \text{Months})_t$	0.804	0.791	0.787	0.833	0.790	-0.014	(-1.12)
Positive Consumption Change $(\pm\ 6\ \text{Months})_t$	0.901	0.895	0.894	0.917	0.897	-0.004	(-0.42)
Negative $\text{Deposits}_{t+1:t+3}$	0.668	0.424	0.296	0.198	0.097	-0.570***	(-46.79)
Negative $\text{Deposits}_{t+1:t+6}$	0.770	0.566	0.435	0.299	0.165	-0.605***	(-49.27)
Negative $\text{Deposits}_{t+1:T}$	0.883	0.754	0.634	0.517	0.309	-0.574***	(-46.53)
Total $\text{Fees}_{t+1:T}$	62.238	55.406	51.428	39.891	19.384	-42.854***	(-19.29)
Average $\text{Fees}_{t+1:T}$	3.350	2.776	2.520	2.040	1.073	-2.277***	(-19.53)
Rating_t	3.956	3.780	3.756	3.645	3.431	-0.525***	(-7.86)
Average $\text{Rating}_{t:T}$	3.255	3.061	3.063	3.101	3.225	-0.030	(-0.57)
$\text{Downgrade}_{t+1:t+3}$	0.000	0.000	0.001	0.000	0.000	0.000	(.)
$\text{Downgrade}_{t+1:t+6}$	0.001	0.003	0.002	0.004	0.001	0.000	(0.00)
$\text{Downgrade}_{t+1:T}$	0.011	0.009	0.012	0.011	0.003	-0.008***	(-3.04)
Observations	2086	2086	2086	2086	2085	4171	

Table 12: Consumption Response of Users with High Deposits over Inflows

This table provides coefficient estimates of OLS regressions estimating the effect of mobile overdraft settings on users' consumption behavior for the top deposits over inflows quintile. *Consumption* is the sum of users' *Card Consumption* and *Cash Withdrawals* in the given month. *Inflows* is the total amount of all incoming transactions a user receives in the given month. *Overdraft Available* is a binary indicator that equals 1 if the user has access to a mobile overdraft in the given month. *Inflows Volatility* refers to the standard deviation of inflows in the year prior to overdraft activation. *Deposits / Inflows* is the ratio of deposits to inflows in the month prior to the overdraft activation. We report t-statistics based on standard errors double-clustered at the NUTS2 and year-month level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable (×100):	$\frac{\text{Consumption}_t}{\text{Inflows}_{t-1}}$				
	(1)	(2)	(3)	(4)	(5)
Overdraft Available$_t$	12.235***	10.648***	10.269***	8.891***	6.501**
	(7.19)	(3.41)	(6.15)	(3.71)	(2.52)
Overdraft Available$_t$ * Inflows Volatility > Median		-3.998		-4.536*	0.511
		(-1.61)		(-1.70)	(0.19)
Overdraft Available$_t$ * Deposits / Inflows > Median			4.010***	4.019	9.916
			(2.82)	(1.04)	(1.29)
Overdraft Available$_t$ * Inflows Volatility > Median * Deposits / Inflows > Median					-10.688
					(-1.20)
Fixed Effects:					
User	Yes	Yes	Yes	Yes	Yes
NUTS3 × Year-Month	Yes	Yes	Yes	Yes	Yes
Standard Error Clusters:					
NUTS2	44	33	44	33	33
Year-Month	42	42	42	42	42
Adjusted R^2	0.256	0.242	0.256	0.242	0.242
User-Year-Month Observations	36,854	5,147	36,854	5,147	5,147

Appendix

A Variable Definitions

Variable	Definition
Age	End of month date minus first day of user's birth year.
Time Since Account Opening	End of month date minus the date when the user completed the account opening procedure.
Female	Indicator variable equal to one if the user is female.
Urban	Indicator variable equal to one if the user lives in a NUTS3 region with a population of at least 500,000 people.
Overdraft Available	Indicator variable equal to one if the user has access to a mobile overdraft in the given month.
Overdraft Enabled	Indicator variable equal to one if the user has access to a mobile overdraft and enabled the credit line by setting a positive user amount in the given month.
Push Notifications Active	Indicator variable equal to one if the user has access to a mobile overdraft, enabled the overdraft, and allowed the mobile app to send notifications to remind her that she uses the overdraft and accrues interest.
Overdraft Amount	Maximum overdraft amount granted to the user by the bank in the given month.
Negative Deposits	Indicator variable equal to one if the user has a negative account balance in the given month.
Overdrawn Amount	Amount of *Negative Deposits*.
Fees	Total amount of fees that the bank deducts from the user's checking account in the given month.
Rating	Consumer credit rating of user i in the given month, ranging from 1 (highest rating) to 6 (lowest rating).
Card Consumption	Total amount of electronic card consumption in the given month.
Cash Withdrawals	Total amount of cash withdrawals in the given month.
Consumption	Sum of *Card Consumption* and *Cash Withdrawals*.
Inflows	Total amount of all incoming transactions a user receives in the given month.
Entertainment	Monthly, electronic consumption expenditures on entertainment.
Shopping	Monthly, electronic consumption expenditures on shopping.
Gastronomy	Monthly, electronic consumption expenditures on gastronomy.
Travel	Monthly, electronic consumption expenditures on travel.
Groceries	Monthly, electronic consumption expenditures on groceries.
Discretionary	Sum of users' monthly expenditures on *Entertainment*, *Shopping*, *Gastronomy*, and *Travel*.
Non-Discretionary	*Card Consumption* minus *Discretionary* spending.
Big Ticket Expense ($> X$ EUR)	Indicator variable equal to one if the user purchased at least one item with a transaction amount exceeding x Euros in the given month.

Internet Appendix to

Perceived Precautionary Savings Motives: Evidence from FinTech

Francesco D'Acunto
Thomas Rauter
Christoph Scheuch
Michael Weber

Table IA1: Effect of Enabling the Overdraft on Users' Consumption Behavior
This table provides coefficient estimates of OLS regressions estimating the effect of enabling the mobile overdraft on users' consumption behavior. *Consumption* is the sum of users' *Card Consumption* and *Cash Withdrawals* in the given month. *Card Consumption* is the user's total amount of electronic card consumption. *Cash Withdrawals* is the user's total amount of cash withrawals from ATMs in the given month. *Inflows* is the total amount of all incoming transactions a user receives in the given month. *Overdraft Enabled* is binary indicator that equals 1 if the user enabled the overdraft in the given month. We report t-statistics based on standard errors double-clustered at the NUTS2 and year-month level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable (×100):	$\frac{\text{Consumption}_t}{\text{Inflows}_{t-1}}$	$\frac{\text{Card Consumption}_t}{\text{Inflows}_{t-1}}$	$\frac{\text{Cash Withdrawals}_t}{\text{Inflows}_{t-1}}$
	(1)	(2)	(3)
Overdraft Enabled$_t$	3.792***	2.512***	0.978***
	(6.40)	(5.91)	(4.60)
Fixed Effects:			
User	Yes	Yes	Yes
NUTS3 × Year-Month	Yes	Yes	Yes
Standard Error Clusters:			
NUTS2	48	48	48
Year-Month	43	43	43
Adjusted R^2	0.254	0.284	0.330
User-Year-Month Observations	517,114	517,149	517,314

Figure IA1: Consumption Pattern around Mobile Overdraft Activation
This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the effect of enabling the mobile overdraft on users' consumption behavior. We estimate model (1) from Table IA1 but replace the *Overdraft Enabled* indicator with separate time dummies, each marking a one-month period (except for event period t-1).

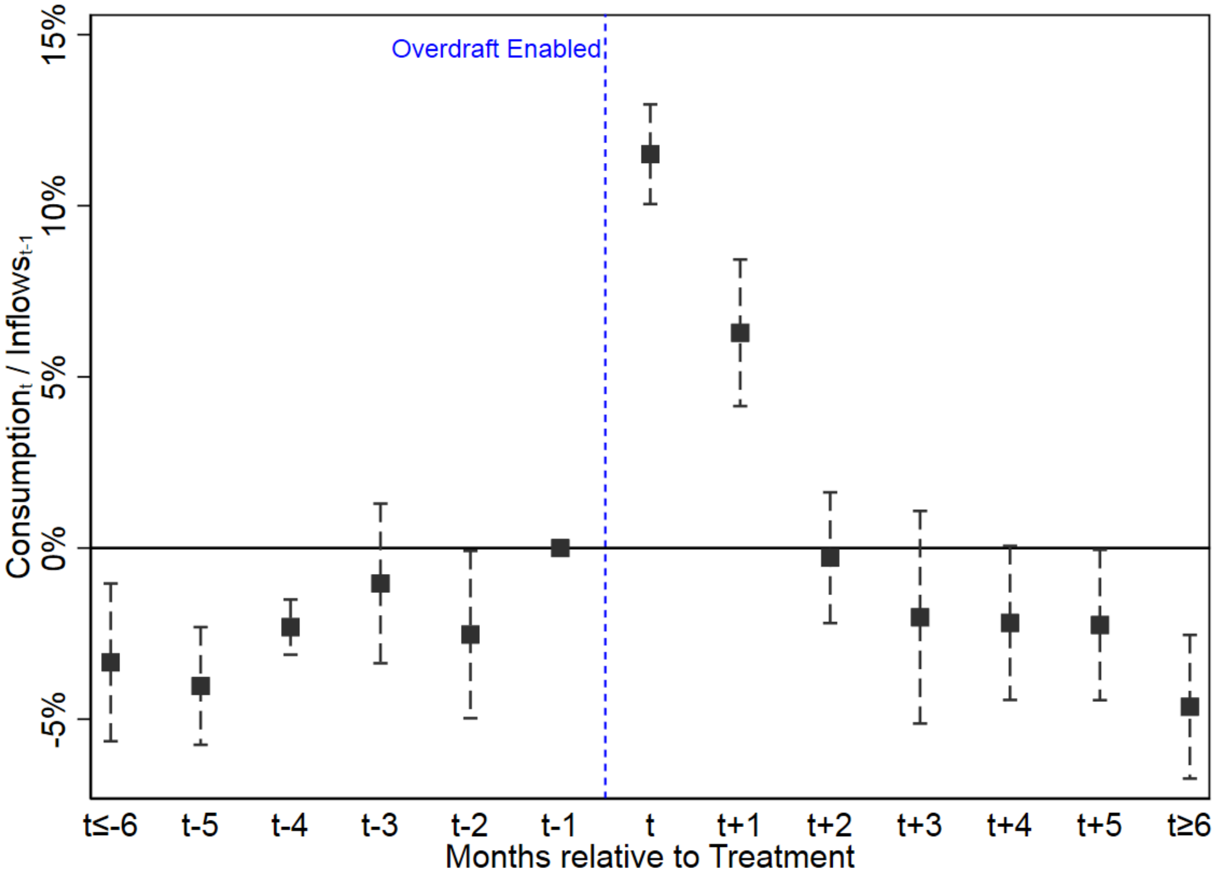

Table IA2: Characteristics of Users in Credit Risk Analysis

This table provides descriptive statistics for users that enter our credit risk analysis. In the first three columns, we report statistics for all users that successfully applied for a mobile overdraft during our sample period. In the next three columns, we restrict our sample to users that experienced at least one rating change after they obtained access to the overdraft. The last two columns test for differences in means across the two groups of users. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

	All Users			Users with Rating Changes			Difference in Means	
	Mean	SD	P50	Mean	SD	P50	Diff. Mean	*t*-Stat.
Age [Years]	33.399	10.260	30.664	32.098	7.615	31.244	1.301	(1.35)
Time Since Account Opening [Years]	0.609	0.525	0.534	0.644	0.490	0.611	-0.035	(-0.56)
Female [0/1=Yes]	0.258	0.437	0.000	0.206	0.408	0.000	0.051	(0.99)
Rating [1-6]	3.719	1.668	4.000	4.175	1.612	5.000	-0.456**	(-2.24)
User Observations	8,015			63			8,078	

Table IA3: Characteristics of Users with and without Mobile Overdrafts

This table reports descriptive statistics for users that enter our "determinants" analysis in Table IA7. Each observation corresponds to one user. *Applies for Overdraft* is an indicator variable equal to one if the user applied for a mobile credit line anytime during our sample period. We measure *Age* and *Time Since Account Opening* at the end of each user's time series. *Female* is an indicator equal to one if the user is a woman. *Inflows* are the average monthly inflows of each user in the sample period. *Inflows Volatility* is the standard deviation of each user's inflows. *Inflows Growth* is the user's average growth of inflows over the sample period. For each variable, we report the number of observations (N), mean, standard deviation (SD), 10% quantile (P10), 25% quantile (P25), median (P50), 75% quantile (P75), and 90% quantile (P90).

	N	Mean	SD	P10	P25	P50	P75	P90
Applies for Overdraft [0/1=Yes]	254,581	0.275	0.447	0.000	0.000	0.000	1.000	1.000
Age [Years]	254,581	35.113	11.135	23.083	26.746	31.997	40.745	51.745
Female [0/1=Yes]	254,581	0.242	0.428	0.000	0.000	0.000	0.000	1.000
Time Since Account Opening [Years]	254,581	1.872	0.860	0.868	1.150	1.807	2.601	3.025
Inflows [Euro]	254,581	792.889	2129.313	0.000	27.838	248.684	881.550	2036.447
Inflows Vola [Number]	241,138	0.000	1.000	-0.866	-0.779	-0.332	0.391	1.398
Inflows Growth [%]	200,202	87.548	220.505	-100.000	-35.041	24.551	115.039	328.333

Table IA4: Descriptive Statistics for Cross-Sectional Consumption Tests

This table reports characteristics of users that enter our cross-sectional consumption analysis is Table 10. *Age* is the user's age at the treatment date in years. *Time Since Account Opening* indicates how long the user has had a checking account at the bank in years. *Inflows Growth* is the user's average growth of inflows 6 months before to 6 months after the overdraft activation. *Deposits / Inflows* is the ratio of deposits to inflows in the month prior to the overdraft activation. For each variable, we report the number of observations (N), mean, standard deviation (SD), 10% quantile (P10), 25% quantile (P25), median (P50), 75% quantile (P75), and 90% quantile (P90).

	N	Mean	SD	P10	P25	P50	P75	P90
Female [0/1=Yes]	13,784	0.213	0.410	0.000	0.000	0.000	0.000	1.000
Age [Years]	13,784	33.055	9.938	22.831	26.160	30.412	37.580	47.830
Time Since Account Opening [Years]	13,784	0.604	0.493	0.068	0.197	0.528	0.862	1.169
Inflows Growth [%]	3,096	249.319	1608.000	2.946	11.664	36.635	110.025	300.058
Deposits / Inflows [%]	10,429	265.528	4470.836	0.353	16.000	56.880	100.900	236.553

Table IA5: Local Continuity Test for User Characteristics around Threshold
This table reports non-parametric estimates for the RD treatment effect of a 250 Euro higher overdraft amount on several user characteristics. The dependent variables are the user's age, gender, and time since account opening at the treatment date as well as the user's level of consumption or inflows in the 3 months prior to the overdraft application. We only use polynomials of order 1 and 2 to avoid overfitting issues (Gelman and Imbens (2018)), apply weights from a triangular kernel because it is the mean squared error (MSE) minimizing choice (Cheng et al., 1997), and employ the MSE-optimal bandwidth selection procedure recommended by Calonico et al. (2014). We report both conventional and robust RD estimates (Calonico et al., 2014, 2018). ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable:	Age	Female	Time Since Account Opening	$\text{Consumption}_{t-3:t-1}$	$\text{Inflows}_{t-3:t-1}$
	(1)	(2)	(3)	(4)	(5)
Conventional	0.905	0.0295	-0.104	-59.48	-202.2
	(0.48)	(0.42)	(-1.21)	(-0.47)	(-0.45)
Robust	1.334	0.0287	-0.145	-72.42	-304.1
	(0.59)	(0.34)	(-1.45)	(-0.47)	(-0.54)
User Observations	1,906	1,906	1,906	1,906	1,906
Order Local Polynomial (p)	1	1	1	1	1
Order Bias (q)	2	2	2	2	2
Bandwidth Left	35.16	39.15	28.56	45.49	47.00
Bandwidth Right	35.16	39.15	28.56	45.49	47.00
Effective Observations Left	271	296	223	377	385
Effective Observations Right	259	276	225	364	371

Figure IA2: Regression Discontinuity Plots for User Characteristics

This figure provides graphical evidence for local continuity in user characteristics at the rounding threshold. We aggregate our data into 16 disjoint bins, calculate the average value, plot this value above the midpoint of each bin, and separately fit 2 linear regressions through all observations on each side of the rounding threshold.

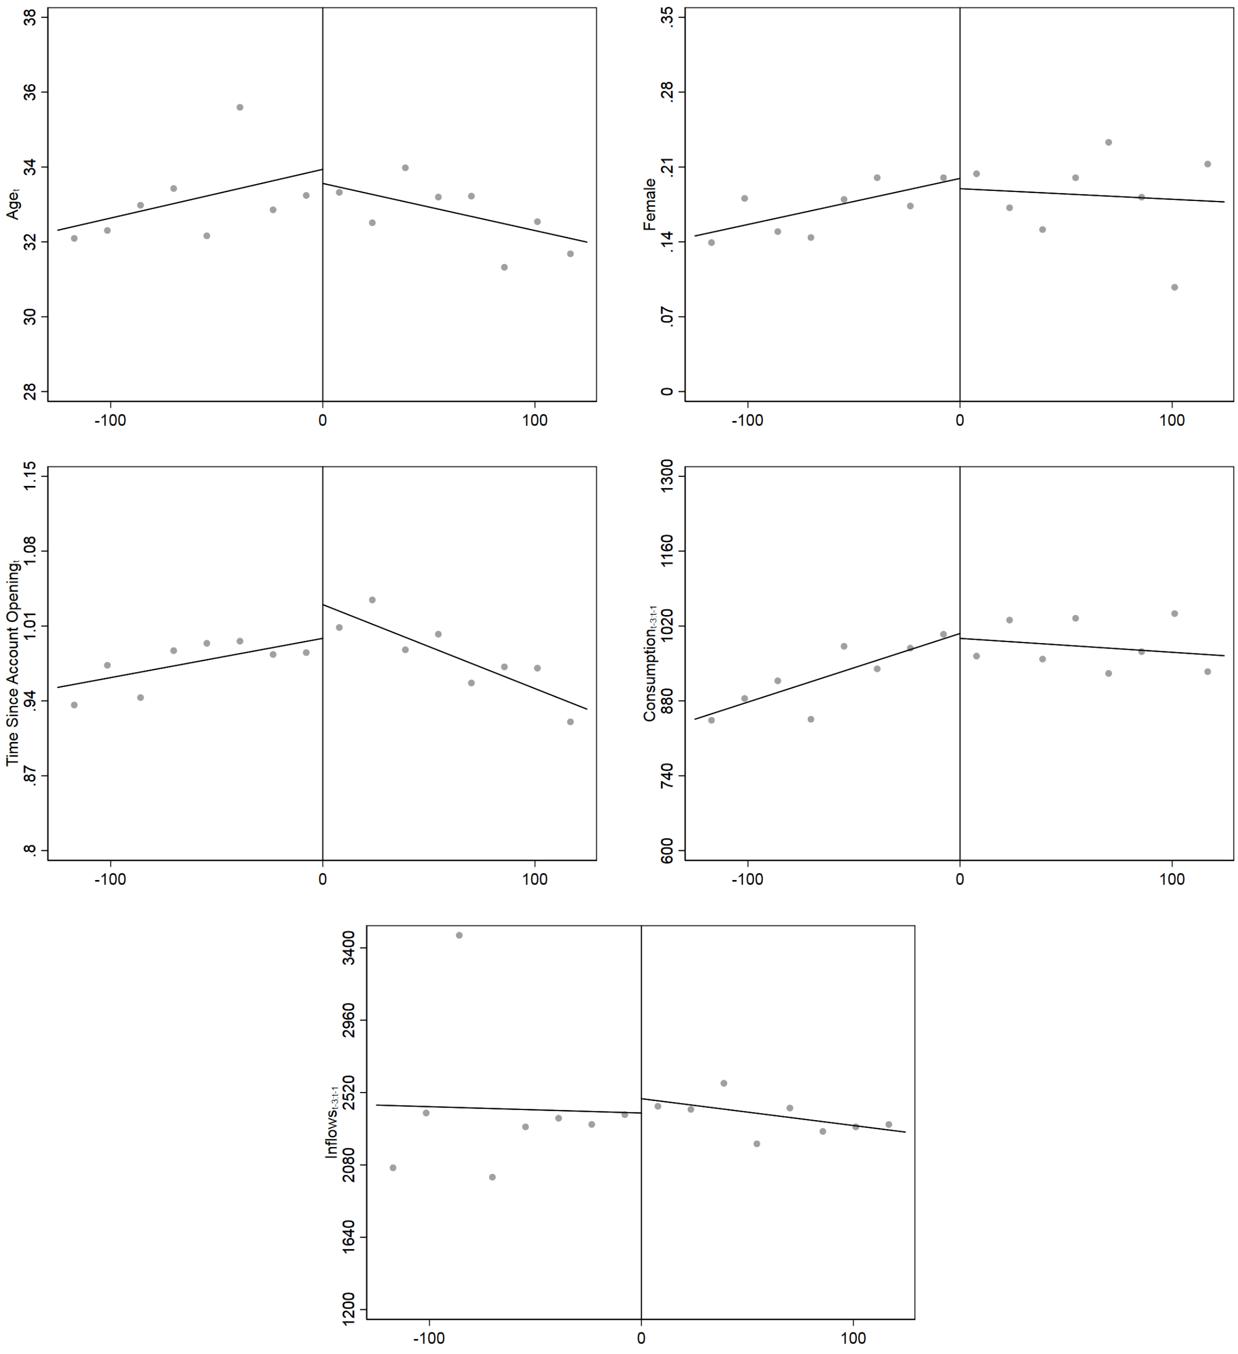

Figure IA3: Robustness Tests for Treatment Manipulation Analysis
This figure reports t-statistics for the treatment manipulation test by Cattaneo et al. (2017) for different polynomial orders and bandwidth choices. The vertical horizontal lines indicate the critical 10% significance levels at which the test rejects the null hypothesis that our running variable is locally continuous around the rounding threshold.

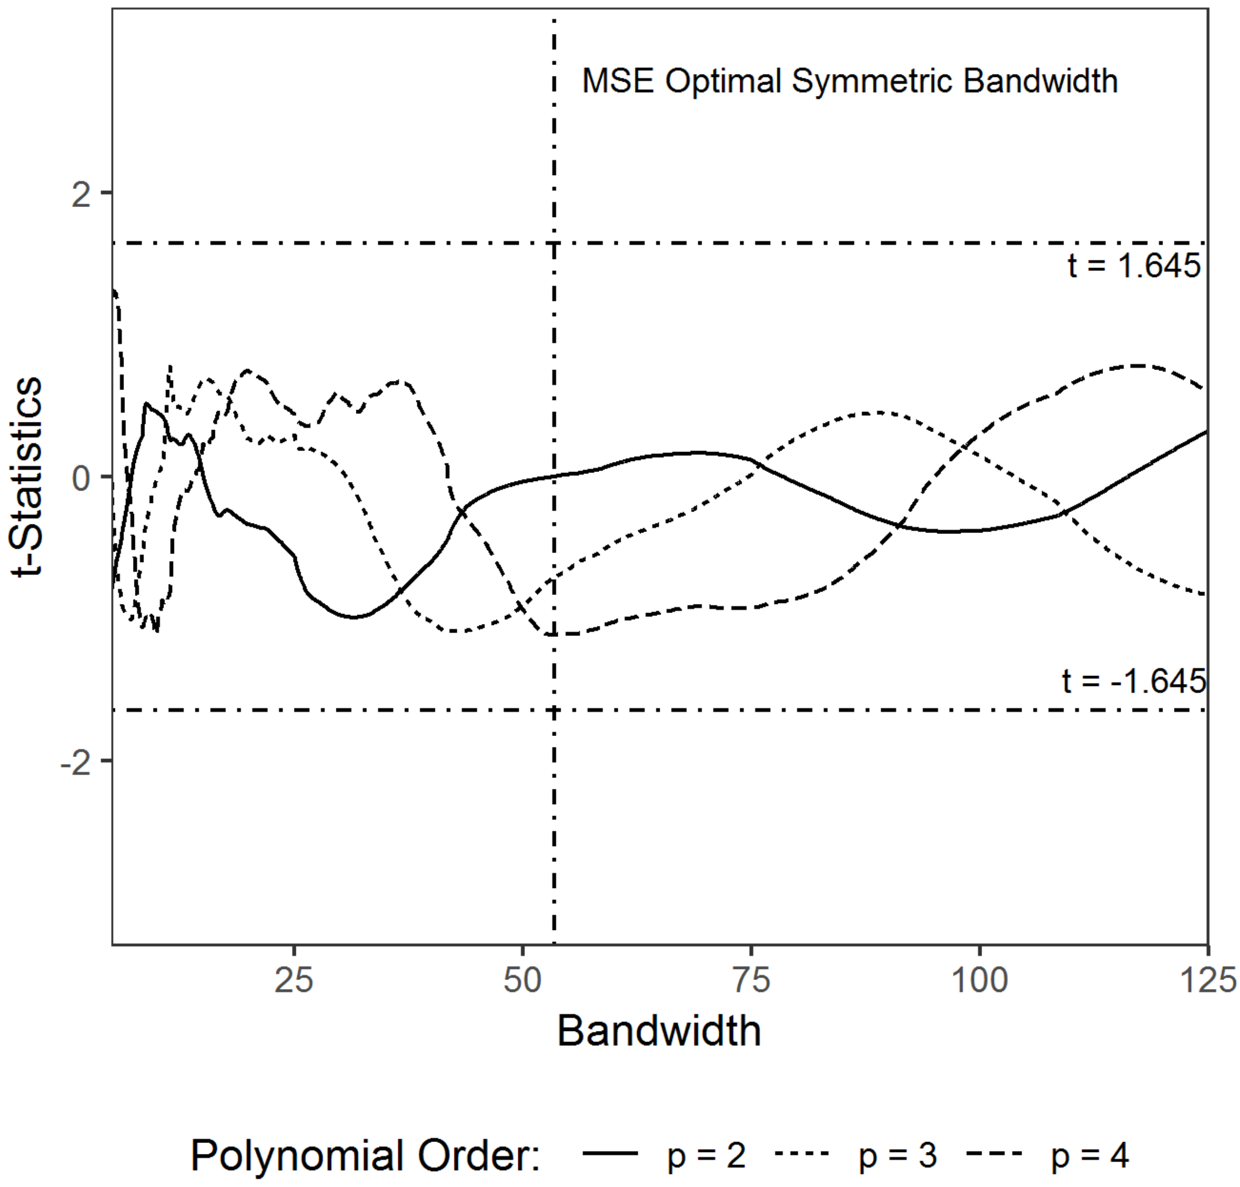

Figure IA4: Sensitivity of RD Consumption Effect to Bandwidth Choice
This figure shows that different bandwidth choices do neither substantially affect the magnitude nor the significance of our main RD consumption effect. Varying the bandwidth is only meaningful over small intervals around the mean-squared-error (MSE) optimal choice (Cattaneo et al., 2020). Bandwidths much larger than the MSE-optimal bandwidth bias the RD estimator, while substantially smaller bandwidths inflate its variance.

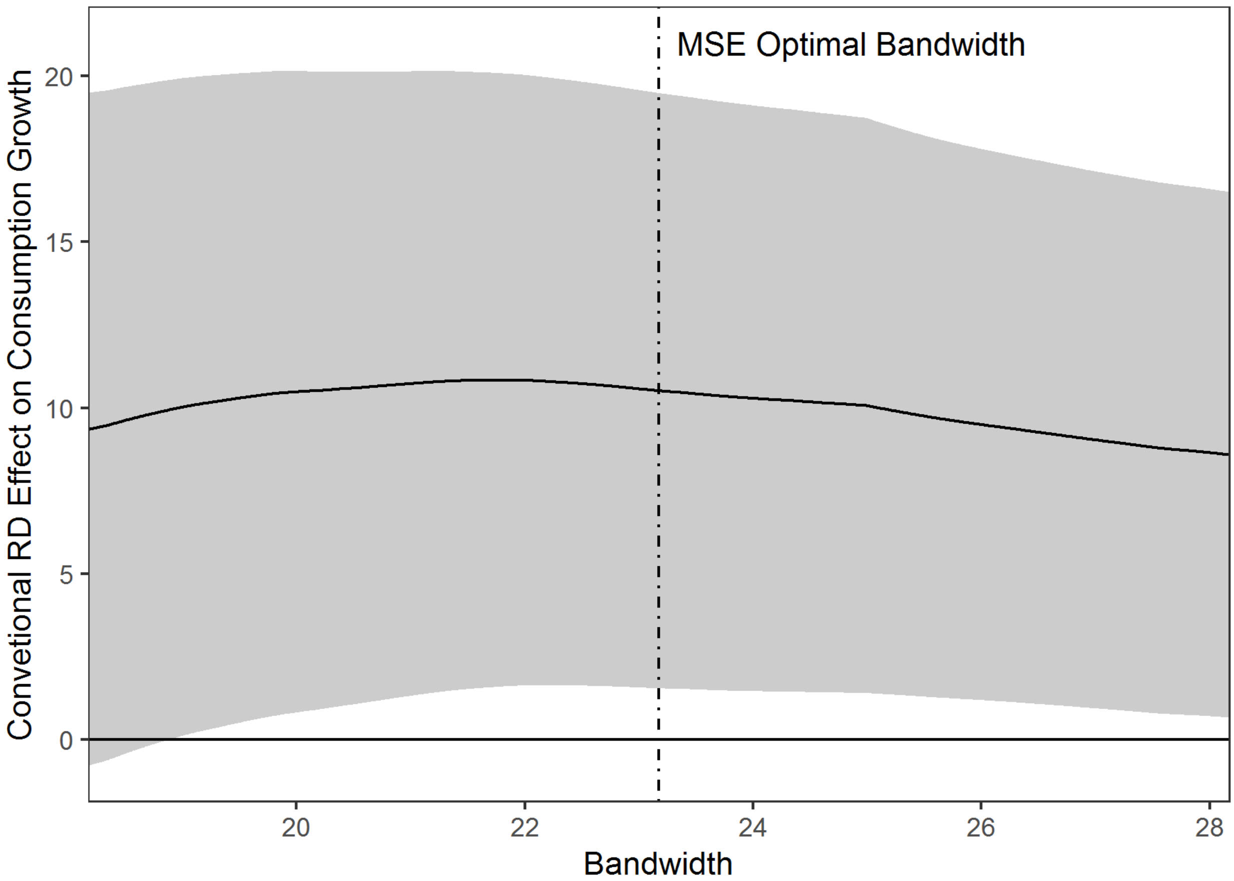

Figure IA5: Sensitivity of RD Results to Observations around Threshold
This figure shows that our main RD consumption effect is robust to excluding data close to the rounding threshold (e.g., Barreca et al., 2011, 2016). We drop users located within the radius $r > 0$ of the rounding cutoff. Specifically, we exclude observations for which $|X_i| \leq r$ (Cattaneo et al., 2020) and illustrate that observations close to the rounding threshold do not drive our results.

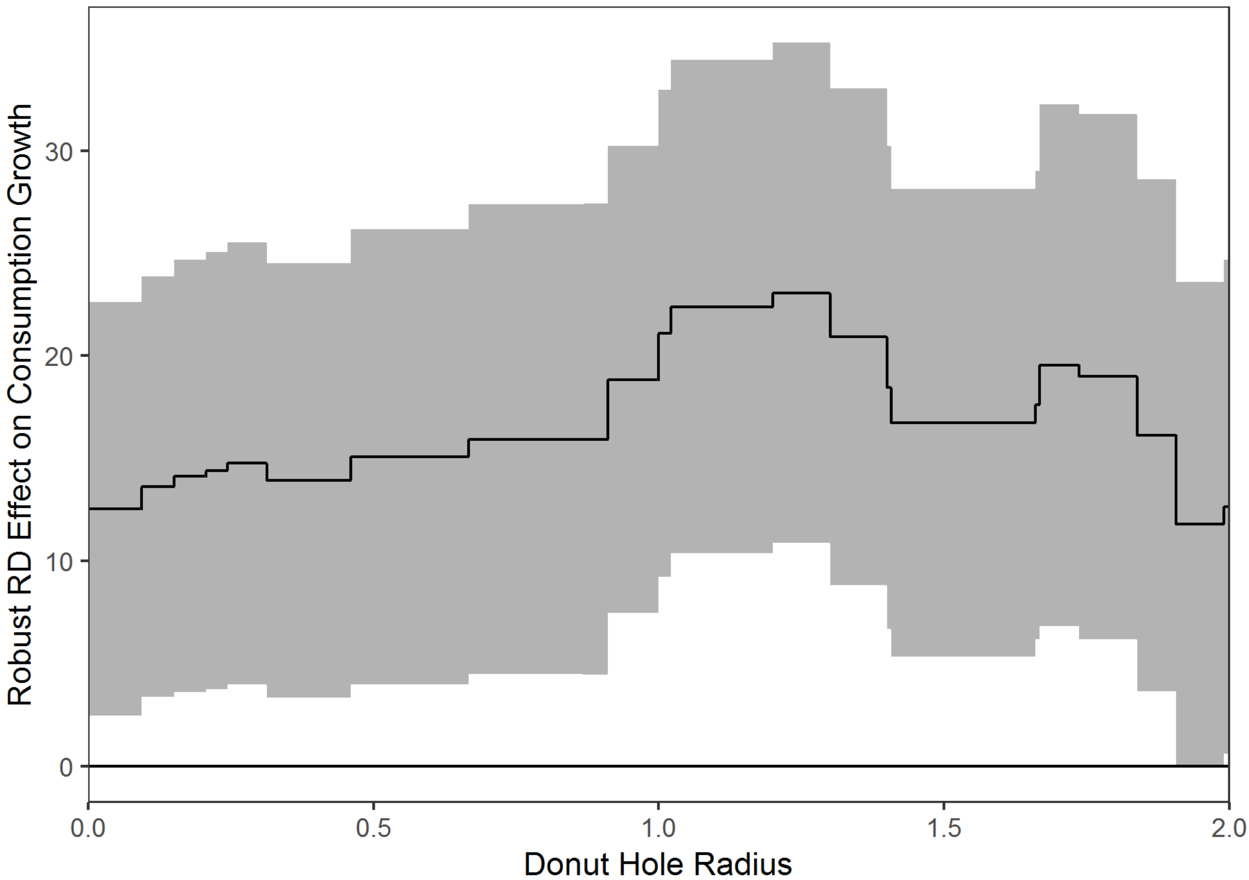

Table IA6: Consumption Growth for Deposits over Inflows Splits
This table presents non-parametric estimates for the RD treatment effect of a 250 Euro higher overdraft amount on users' consumption behavior. The dependent variable is the difference in average consumption three months before and after the treatment, normalized by the average account inflows three months prior to the overdraft application. We residualize users' consumption growth rate with country $\times$ year-month fixed effects to ensure that we compare treated and control users from the same European country at a similar point in time. We only use polynomials of order 1 and 2 to avoid overfitting issues (Gelman and Imbens (2018)), apply weights from a triangular kernel because it is the mean squared error (MSE) minimizing choice (Cheng et al., 1997), and employ the MSE-optimal bandwidth selection procedure recommended by Calonico et al. (2014). We report both conventional and robust RD estimates (Calonico et al., 2014, 2018). We control for *Age*, gender (*Female*), and *Time Since Account Opening* in all specifications. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

	$\frac{\text{Deposits}_{t-1}}{\text{Inflows}_{t-1}}$ < Median		$\frac{\text{Deposits}_{t-1}}{\text{Inflows}_{t-1}}$ > Median	
	(1)	(2)	(3)	(4)
Conventional	12.22*	14.04*	8.792	10.84
	(1.76)	(1.79)	(1.56)	(1.59)
Robust	14.51*	15.98*	11.27*	13.09*
	(1.87)	(1.88)	(1.67)	(1.69)
Covariates	Yes	Yes	Yes	Yes
Observations	932	932	935	935
Order Local Polynomial (p)	1	2	1	2
Order Bias (q)	2	3	2	3
Bandwidth Left	20.96	35.35	29.16	43.49
Bandwidth Right	20.96	35.35	29.16	43.49
Effective Observations Left	78	135	110	180
Effective Observations Right	69	123	115	176

Table IA7: Determinants of Mobile Overdraft Application

This table provides coefficient estimates of OLS regressions estimating the determinants of users' decision to apply for a mobile overdraft. Each observation corresponds to one user. *Applies for Overdraft* is an indicator variable equal to one if the user applied for a mobile credit line anytime during the sample period. We measure *Age* and *Time Since Account Opening* at the end of each user's time series. *Female* is an indicator equal to one if the user is a woman. *Rating* is the user's average rating over the sample period. *Inflows* is the user's average inflows, *Inflows Volatility* is the standard deviation of each user's inflows, and *Inflows Growth* is the user's average inflow growth over the sample period. We report t-statistics based on standard errors clustered at the NUTS2 in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Dependent Variable:	Applies for Overdraft			
	(1)	(2)	(3)	(4)
Age	0.000	0.000	0.000	-0.000
	(1.25)	(0.88)	(0.03)	(-0.37)
Female	-0.022***	-0.021***	-0.020***	-0.022***
	(-5.94)	(-5.59)	(-5.45)	(-6.16)
Time Since Account Opening	-0.045***	-0.044***	-0.045***	-0.042***
	(-3.38)	(-3.34)	(-3.32)	(-3.12)
Log(Inflows)		0.011***	-0.014***	-0.014***
		(6.60)	(-5.13)	(-5.20)
Inflows Volatility			0.025***	0.026***
			(9.98)	(9.95)
Inflows Growth				-0.002
				(-0.53)
Fixed Effects:				
NUTS3	Yes	Yes	Yes	Yes
Standard Error Clusters:				
NUTS2	48	48	48	48
Adjusted R^2	0.041	0.042	0.041	0.039
User Observations	130,597	130,570	117,539	107,372