ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

D'Acunto, Francesco; Rauter, Thomas; Scheuch, Christoph; Weber, Michael; Weber, Michael

Working Paper

Perceived Precautionary Savings Motives: Evidence from FinTech

CESifo Working Paper, No. 8123

Provided in Cooperation with:
Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: D'Acunto, Francesco; Rauter, Thomas; Scheuch, Christoph; Weber, Michael; Weber, Michael (2020) : Perceived Precautionary Savings Motives: Evidence from FinTech, CESifo Working Paper, No. 8123, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:
https://hdl.handle.net/10419/216519

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

CESifo WORKING PAPERS

8123
2020

February 2020

Perceived Precautionary Savings Motives: Evidence from FinTech

Francesco D'Acunto, Thomas Rauter, Christoph Scheuch, Michael Weber

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 8123*

Perceived Precautionary Savings Motives: Evidence from FinTech

Abstract

We study the spending response of first-time borrowers to an overdraft facility and elicit their preferences, beliefs, and motives through a FinTech application. Users increase their spending permanently, lower their savings rate, and reallocate spending from non-discretionary to discretionary goods. Interestingly, liquid users react more than others but do not tap into negative deposits. The credit line acts as a form of insurance. These results are not fully consistent with models of financial constraints, buffer stock models, or present-bias preferences. We label this channel perceived precautionary savings motives: Liquid users behave as if they faced strong precautionary savings motives even though no observables, including elicited preferences and beliefs, suggest they should.

JEL-Codes: D140, E210, E510, G210.

Keywords: household finance, consumption, behavioral finance.

Francesco D'Acunto
Boston College Carroll School of Management
Chestnut Hill, MA 02467, USA
dacunto@bc.edu

Thomas Rauter
University of Chicago Booth School of Business, Chicago, IL 60637, USA
thomas.rauter@chicagobooth.edu

Christoph Scheuch
Vienna Graduate School of Finance & WU, 1020 Vienna, Austria
christoph.scheuch@wu.ac.at

Michael Weber
University of Chicago Booth School of Business, Chicago, IL 60637, USA
michael.weber@chicagobooth.edu

February 12, 2020
We thank Fabian Nagel, Tom Kim, Hannah Amann, and especially Federica Ansbacher for excellent research assistance. We appreciate helpful comments from Sumit Agarwal, Greg Buchak, Indraneel Chakraborty, Francisco Gomes, Sabrina Howell, David Low, Amit Seru, Wenlan Qian, and workshop participants at the Columbia "New Technologies in Finance" Conference, the 2019 Red Rock Finance Conference, the 2019 Summer Finance Conference at ISB, the 7th ABFER Annual Conference, the 2019 LBS Summer Finance Symposium, the 2019 CFPB Conference, and the 2020 AEA and AFA Annual Meeting. 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

During recessions, monetary and fiscal policy aim to stimulate consumption through credit, because household consumption comprises the largest share of gross domestic product (Agarwal et al., 2017). At the same time, household 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; Mian and Sufi, 2015; Mian et al., 2017). Understanding the micro-level channels through which household credit affects consumption and saving decisions is thus important to understand the dynamics of business cycles and to inform the design of effective credit-based expansionary policies.

In this paper, we introduce a unique FinTech setting in which we observe the extensive margin of credit—initiation of overdraft facilities to first-time borrowers—combined with the possibility to elicit *directly* a set of preferences, beliefs, and individual perceptions that theoretical and empirical research commonly relates to households' saving and consumption behavior. These dimensions, which are typically unobserved in large-scale micro data, include risk preferences, patience, expectations about future employment status, future large expenses, medical expenses, bequest motives, subjective life expectancy, as well as financial literacy and generalized trust in others. We use this setting to estimate the consumption effects of providing first-time borrowers with a credit facility and to test directly for a broad set of typically unobserved channels that, based on earlier theoretical and empirical research, might explain the effects of credit on spending.

In our baseline analysis, we estimate a set of double-differences specifications in which we compare bank users' change in consumption spending after activating the overdraft facility relative to before and relative to users whose overdraft facilities are not yet activated. The average user increases her spending by 4.5 percentage points of income inflows after credit is available relative to before and relative to users whose overdraft facility is not yet active. This positive effect on spending does not revert fully over time—we detect permanent increases in spending and drops in savings for the average user.

The surprising result is that the baseline effect is heterogeneous across users, but not based on dimensions earlier research proposed, such as motives to smooth consumption or liquidity constraints. Instead, Figure 1 shows the effect increases monotonically in the liquidity of users. The higher the ratio of liquid deposits to income inflows, the stronger is the spending reaction to the provision of credit. Moreover, the higher is this ratio, the more permanent is the effect as we document below.

Liquid users seem to behave as if they have strong precautionary savings motives

Figure 1: Consumption Response by Deposit-to-Income

Notes: 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 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.

and hence accumulate liquid savings before the overdraft facility is available. They have inflows over the previous twelve months similar to other users but save a larger fraction of their income by spending substantially less. Once the overdraft becomes available, consistent with an insurance effect of the facility, liquid users start spending more of their existing liquid savings and decrease their savings rate permanently. Interestingly, liquid users do *not* use the overdraft facility, in the sense that only 10% of them ever tap into negative deposits and make use of the credit line. The availability of an overdraft facility makes these users dissave part of their existing liquidity but not more than that, again consistent with strong precautionary savings motives that are reduced once the facility is available.

The spending response by liquidity is a robust feature of the data. First, we rule out the amount or volatility of inflows into the account before the credit line is available drives our results. Moreover, the results are similar if we only consider individuals who consume large parts of their inflows, for alternative definitions of deposits-to-inflows, for users in areas that have high versus low savings rates—and hence potentially different social norms about the importance of savings—for urban versus rural users, or users living in former communist countries and others. Thus, peculiar norms of saving or aversion to

debt in specific European countries, such as Germany, do not drive our results.[1]

Our results are puzzling to the extent the previous literature typically found that liquidity constraints drive borrowers' reaction to the availability of credit. Hence, the least liquid users should react the most in terms of spending—the opposite of our finding. For instance, research that studied increases in the intensive margin of credit, such as increasing existing credit card limits, finds users at the constraint (i.e., that maxed out their previous limit) react more to the increase in credit lines (see, e.g., Aydin, 2015; Agarwal et al., 2017; Gross and Souleles, 2002). The crucial difference of our setting is that we study the provision of credit lines to borrowers who did not have access to credit beforehand (extensive margin of credit).[2] The insurance role of obtaining a credit facility for the first time does not apply to borrowers who already had access to credit, and hence the effect we document could not arise in earlier settings that vary the intensive margin of credit. Another difference between our results and earlier research is that our liquid users do not tap into negative deposits, but only consume out of their existing savings and new inflows, which suggests that changes in credit limits or in the intensive margin of credit would barely have any effects on them.

To understand whether standard determinants of precautionary savings motives explain our results, we exploit the fact that our FinTech provider allows us to reach users directly through their app and elicit their preferences, beliefs, and motives. First, risk aversion and/or prudence are important potential drivers of savings behavior independent of income volatility (Gomes and Michaelides, 2005). We detect no economic or statistical differences in the levels of risk aversion and patience across users based on their liquidity. Moreover, users who have high subjective beliefs about future large expenses, such as medical expenses, or about future income uncertainty (Guiso et al., 1992; Ben-David et al., 2018) might save a larger fraction of their income. Even here, we find no systematic differences in the cross section of deposit-to-inflows. We do not find systematic differences in terms of subjective life expectancy, beliefs about future employment status, health conditions, or bequest motives, either. Finally, we do not find any systematic differences in terms of financial literacy or generalized trust—liquid users understand the incentives to borrow and save similarly to other users.

Ultimately, the majority of users—all those above the bottom 40% by liquidity—behave as if they had precautionary savings motives before accessing the credit line.

[1] We thank Amit Seru for suggesting sample splits based on areas with different norms about saving and borrowing within and across European countries.

[2] We discuss the institutional setting in detail in Section 2.

These motives disappeared with the activation of credit lines, which we could interpret as a form of insurance against future negative states. At the same time, none of the standard explanations for precautionary savings motives appear to have the potential to explain our results, either individually or jointly. For this reason, we label our mechanism "perceived precautionary savings"—users perceive strong precautionary savings motives, but neither their preferences, subjective or objective beliefs, or other characteristics we elicit directly based on earlier theoretical and empirical research, explain these motives.

We argue that models of buffer stock consumption (Deaton, 1991; Carroll, 1997), buffer stock consumption with durables (Guerrieri and Lorenzoni, 2017; Aydin, 2015), the standard 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. Non-standard preferences such as present-biased preferences are also unlikely to explain all our results. If the users who 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 is not true in the data. 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 response to income payments. Ultimately, our results call for further theoretical and empirical evidence on potential alternative drivers of precautionary savings motives, either based on standard or non-standard beliefs. For instance, a promising direction could be to allow heterogeneity in the extent to which agents have biased beliefs regarding the occurrence of small-probability events, such as catastrophic states of the world.

A remaining concern with our baseline double-differences analysis is that users might activate the overdraft facility endogenously when they know they are about to make a large expense. The facts that the spending effect for liquid households is permanent and that elicited expectations about upcoming large expenses do not differ across liquid and illiquid users reduce this concern.

To address this concern directly, we propose two tests. First, we repeat our double-differences analysis restricting the sample only to users who activated their overdraft facility in the first month in which the bank made it available.[3] The rationale is that

[3]We thank Greg Buchak and Michaela Pagel for suggesting this test.

the timing at which the bank decided to make the facility available can barely coincide with the timing in which users would have endogenously decided to activate the facility, because users did not know that the overdraft facility would have been made available at that time. Although the sample size drops substantially in this test, we estimate the same monotonic pattern in terms of change in spending by liquidity quintiles and we can reject the null that the changes are zero at conventional levels of significance for the top two quintiles.

For the second test, we propose a sharp regression-discontinuity design (RDD) that exploits the rule the FinTech provider uses to compute the 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 app allows users to open the approved line of credit automatically in just a few seconds—and hence there is no scope for manipulating one's assignment to the high- or low-treatment group.

In this RDD, we compare users that are observationally indistinguishable, as we document directly, but end up being assigned different overdraft limits based on small differences in their income inflows. This assignment combines the extensive margin of credit, so that the insurance effect of the facility exists for everybody, with a higher or lower intensive margin in terms of credit limit. The RDD confirms the consumption response to the availability of credit. We cannot use this RDD for our baseline analysis and heterogeneity tests, because by construction the RDD only uses a small subset of users who are close to the sharp threshold. Heterogeneity in the RDD sample is minimal by construction given that the treated and control users are almost identical in most characteristics.

Ultimately, our results might provide novel insights about the effects of providing credit to households over the business cycle. Providing insurance against potential negative future spending shocks appears to increase the spending of households with high deposit-to-income ratios, but this policy intervention would not increase their debt positions or interest payments, because these users would not draw on the overdraft facilities. If anything, providing insurance to these households at times of economic slumps might increase aggregate demand swiftly, because households that hoard cash due to perceived precautionary savings motives might end up spending parts of their savings. Also, providing insurance to perceived precautionary savers might be virtually costless based on our results, because these individuals would not end up paying any interest and would not accumulate debt over time.

The main challenge to policies based on our findings might be political in nature. During times of economic crisis, policy-makers would provide virtually costless insurance to *wealthier* households (in terms of liquid wealth) to nudge them to spend more instead of providing subsidies to induce spending by all households, including those who might become risky borrowers over time. Whether governments might find enough political support for this type of policy is a matter for future research in political economy and political sciences.

Methodologically, the paper suggests a compelling reason for macroeconomists and financial economists to use FinTech settings for research on household borrowing, saving, and spending. FinTech platforms allow the researcher to access users directly in a logisitically simple way that barely involves any costs, because users can be contacted interactively through apps. Easy direct access to users allows the elicitation of typically unobserved characteristics, as in this paper, and has the potential for field experiments to gauge the causal effects of theoretical channels that act through preferences, beliefs, and perceptions on actual high-stake decision-making.

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. 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 identification 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. Contrary to US banks, our FinTech lender does not offer credit cards, which is common in the countries in which it operates. Moreover, note EU legislation does *not* require banks to provide any credit facilities to their customers, which is the reason why our setting allows to test for the effects of providing an overdraft facility to users for the first time that allows borrowing for consumption spending.[4]

[4]For instance, for the case of Germany, the largest European country, which enters our sample, see the information sheet by the national bank regulator, BaFin: https://www.bafin.de/DE/Verbraucher/Bank/Produkte/Dispokredit/dispokredit_node.html.

In terms of day-to-day operations, 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.

In February 2015, the bank started to offer a mobile credit line to users across several European countries. Residents of these countries with a sufficiently high credit score are eligible for the mobile overdraft facility. Customers can activate the credit line directly in their mobile app within one minute and receive a maximum overdraft amount that ranges 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 6, we describe the bank's loan granting and credit allocation process in detail. Users who 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 granted to them by the bank. Users pay an annual interest rate of approximately 10 percentage points on their used overdraft amount, which the bank charges once per 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, receive unemployment benefits, or experience direct debit reversals.

3 Data and Descriptive Statistics

3.1 Data Sources and Sample Selection

Our sample consists of users who received an overdraft between February 2015 and September 2017. 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 directly. 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 bank 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. In a set of robustness analysis, we show that our results do not change if we use other thresholds or alternative filters to screen out users who are unlikely to be main account users.

For all the users in our sample, we first obtain detailed transaction-level consumption data, credit line information, and personal user characteristics from the bank. Specifically, we obtain information about the type, amount, and timestamp of all financial transactions that pass through users' checking accounts between February 2015 and March 2019. To protect users' identity, 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. Our raw dataset contains 58,310,004 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 opened the account, and ends with the closing of the account or the last month of our sample period (March 2019). We code observations of our flow variables as zero if the user did not have any corresponding financial transaction in the given month.

An important feature of the transaction-level data is the classification of financial transactions into categories, which the bank provides to us. The first classification consists of six broad categories based on the transaction method: (*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. In a second classification, which captures the scope of each transaction, the bank categorizes each debit-card payment into one of seventeen merchant category code (MCC) groups, based on the vendor that receives the payment. MCC groups specify the vendor's industry and allow us to identify which type of good or service the account holder purchased. The seventeen MCC groups cover a broad range of users' consumption spending and include both discretionary (e.g., entertainment, shopping, or gastronomy) and non-discretionary consumption categories (e.g., groceries, family, or utilities/furniture).

In addition to the transaction-level data, we obtain data on the activation and usage of credit lines. This dataset contains granular information about the application date, the granted overdraft amount, and activating users' financial characteristics. 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-line dataset contains the complete history of all overdraft setting changes that users made once after activating the credit facility.

Finally, the bank provides us with additional data on the demographic and personal information of 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 account holders.

3.2 Descriptive Statistics

Table 1 provides descriptive statistics for our main overdraft sample. We define all the variables we use in the empirical analysis in Appendix Table A1. 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.[5]

Our final dataset contains 603,157 user-month observations for 36,005 unique individuals who obtained a mobile credit line between February 2015 and September 2017. As is common in FinTech settings,[6] the user base of the 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 2,220 Euros, and opened the mobile bank account 1.6 years ago. Female users represent about one fifth (21%) of the sample, and about half of the users live in large cities with more than 500,000 inhabitants (53%).

In terms of spending, all overdraft users combined spent a total of 441 million Euros via their mobile checking account over our sample period. On average, these individuals consume 48% of their 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 61 cents of discretionary items. Main account users have access to the bank's mobile overdraft facility in 91% of all user-months. The average maximum overdraft amount equals 1,561 Euros.

[5] 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).

[6] For instance, see D'Acunto et al. (2019a), Gargano and Rossi (2018), D'Acunto et al. (2019b), and D'Acunto and Rossi (2020).

4 The Effect of Mobile Overdrafts on Users' Spending Behavior

In this section, we first assess how the average user changes her spending behavior once she activates the overdraft facility on her mobile phone. We then move on to document a set of heterogeneity results in users' response to the availability of the credit line. We find that users with higher ratios of deposits to income flows are the ones reacting the most in terms of increasing their spending after they obtain access to the overdraft facility.

Our baseline analysis considers individuals' monthly consumption expenditures and uses a double-differences (DD) design. The DD estimator compares changes in the level of spending around the activation of mobile overdrafts between individuals that already have and those that do not have access to the credit line yet. We restrict the analysis to mobile overdraft users only to address the endogeneity of being eligible for the overdraft facility, which by construction depends on a set of demographic characteristics that might be correlated with changes in spending patterns over time. Moreover, the DD design allows us to address the concern that individuals who never activate the overdraft despite being eligible might differ systematically from active users based on unobserved characteristics. These characteristics might also predict spending patterns over time, such as users' inattention to their financial situation or the possibility that some individuals use the bank account only for the purpose of depositing their money but do not use it for regular consumption activities.

The baseline DD design does not account for the possibility that unobserved shocks at the individual level determine the *timing* of users' overdraft activation and spending patterns over time. We tackle this concern directly later in the paper in two ways. First, we repeat the analysis after restricting the sample to users who activated the overdraft facility in the first month in which the bank made it available. The concern of endogenous timing of activation is less compelling in this subsample, because these users could not foresee this option well in advance before activation. Second, we propose a sharp RDD that only compares users at the time they activate their facilities, relative to before.

Despite the potential endogeneity of the timing of activation, we use the DD design for the baseline analysis, because this design allows us to assess potential channels driving the consumption response in the full sample. Instead, by construction, the RDD is based on estimating local treatment effects among a restricted subsample of users in the overall population, based on their closeness to the regression discontinuity we exploit.

4.1 Overall Consumption Expenditure Response

For our DD analysis, we estimate the following OLS regression model:

$$\text{Consumption Spending}_{i,t} = \beta \times \text{Overdraft Available}_{i,t} + \text{Fixed Effects}_{i,t} + \varepsilon_{i,t}. \quad (1)$$

The dependent variable is the sum of all cash withdrawals and debit-card 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 that equals 1 if the user has access to a mobile overdraft facility in the given month.

We include user fixed effects to absorb time-invariant systematic differences in consumption spending patterns across overdraft users, such as differences in occupation, gender, cultural background, or education. We also include NUTS3×year-month fixed effects to account for concurrent but unrelated time-varying economic or institutional changes within local, sub-national districts, called NUTS3 in the European Union, which correspond to US counties.

We draw statistical inferences based on double clustered standard errors at the NUTS2 (region) and year-month levels because consumption patterns are likely correlated cross-sectionally and over time within closeby locations. We use a broader geographic definition for the spatial clustering of standard errors, because the NUTS2 level, which corresponds roughly to the state level in the US, fully subsumes the NUTS3 level in the EU's administrative classification. This level of clustering is more conservative than NUTS3, because it allows for correlation of unknown form across the residuals not only across users in the same county, but also across users in neighboring counties that belong to the same region. Because cross-sectional residuals of decision units are typically positively correlated across space, the NUTS 2 level of clustering is also more conservative than clustering standard errors at the individual level. We verify this claim directly by repeating our baseline estimations separately using different levels of clustering, including different geographic partitions and individual-level clustering, and indeed we find that the confidence intervals around our point estimates are wider when we cluster at the NUTS 2 level than at the individual level (see the bottom right panel of Figure 6).

In Table 2, we present the estimated effects of mobile overdraft availability on consumption behavior. Column (1) documents an average positive effect of mobile overdrafts on users' monthly consumption expenditures. In terms of magnitude, overall cash and card-based consumption normalized by the user's lagged account inflows increases by

4.572 percentage points (t-statistic: 11.11), which corresponds to an increase of approximately 9.5% relative to the pre-treatment sample mean (4.572/48.11). In Columns (2) and (3), we differentiate between cash withdrawals, i.e. cash spending, and debit-card spending. We find that account holders increase their spending significantly through both payment types. In relative terms, the increase in card-based spending (coefficient: 3.152; t-statistic: 9.71) is larger than the increase in cash withdrawals (coefficient: 1.345; t-statistic: 7.73), and accounts for approximately 70% of users' overall consumption response. Finally, in Column (4), we study the change in the ratio between discretionary and non-discretionary spending and find that users increase their discretionary spending by 2.9% relative to their non-discretionary spending.

Apart from users' timing to activate the overdraft, we can interpret our DD estimates causally only subject to the assumption that the spending behavior of treated and control users would have followed parallel trends in the absence of the activation of the credit line. This parallel-trends assumption is untestable, but we can directly assess whether the spending behavior of treated and control users followed parallel trends before the shock. In Figure 2, we provide graphical evidence that treated and control users have parallel and almost identical consumption patterns during the time period leading up to the mobile overdraft activation. Treated users sharply increase their spending during the first two months in which they have access to the credit line. Subsequently, the treatment effect stabilizes at around 2% higher consumption relative to inflows and does not fully revert in the long run. We can reject the null of no differences between the spending of treated and control users throughout the sample period and at any horizons we observe in the data after activation. As we discuss below, these baseline effects mask substantial heterogeneity—some groups of users face a permanent spending effect of the overdraft facility, whereas for other users the effect dissipates over time.

5 Heterogeneity and Perceived Precautionary Savings

Our results so far have focused on average effects across users. To investigate the economic channels behind the baseline results and to assess the extent to which existing models might explain them, we study the heterogeneity of the baseline consumption effect in the cross section of users (Jappelli and Pistaferri, 2017).

The observational data we have do not include direct information on a set of char-

acteristics that are crucial to assess the relevance of different models, such as risk aversion, bequests motives, and beliefs. This lack of information typically makes it hard for empirical studies to pin down channels. To solve this problem, we exploit the unique feature that our FinTech bank can communicate with users through the online app to elicit preferences and beliefs *directly* through a survey instrument. This elicitation is a methodological contribution of our paper and the main reason for why the FinTech nature of the lender with which we cooperate is central to our analysis. Without the FinTech app interface, accessing traditional bank users to elicit characteristics directly would require either inviting these individuals to a human-subjects laboratory or sending surveys via mail or phone call, which would limit substantially the scope and possibility of the intervention. We see our paper as one of the very first papers that emphasize the potential for using app-based online bank settings to elicit characteristics of borrowers that would otherwise be barely observable in most real-time studies of borrowing and spending.

5.1 Consumption Smoothing? Income Growth and Age

We start by assessing the heterogeneity of the baseline effect across two characteristics we can observe in the field data, that is, (*i*) the growth of income inflows 6 months after overdraft activation relative to 6 months before activation; and (*ii*) users' age.

The permanent income hypothesis (PIH) suggests agents want to smooth their consumption over time. Empirically, income paths are increasing early in life before flattening out. Hence, the PIH predicts younger users and users with a steeply increasing income path should be more likely to use the overdraft facility to smooth their consumption and borrow against their higher future income to increase their spending on impact relative to other users.[7]

To assess this prediction, we split the sample into quintiles based on the two observable characteristics listed above. We then repeat the baseline DD analysis of Equation (1) adding a set of interactions between an indicator for whether the user belongs to the quintile of the respective characteristic and our overdraft treatment variable. To make the results easier to visualize, in Figure 3 we report the estimated coefficients and 95% confidence intervals for each interaction term in graphical form.

The top panel reports the effects across quintiles of income growth. The estimated

[7]Not only the level, but also the uncertainty of the future income path might be relevant to predict spending (Guiso et al., 1992). We tackle this point directly below using past income volatility as well as with a survey instrument to elicit uncertainty about future income flows.

effect does not exhibit any systematic difference for users in the bottom two quintile relative to users in the top two quintiles, and we do not reject the null hypothesis that the effect is equal across all of the quintiles.

In the bottom panel of Figure 3, 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 and 45 years. We can plausibly argue that users across these quintiles are on different consumption life-cycle paths. The bottom panel of Figure 3, though, shows vividly the lack of heterogeneity of the baseline effect across the age distribution. In terms of magnitudes, we estimate coefficients that range between 3% and 4% in each quintile. Moreover, we cannot reject the null that the effects are the same across all quintiles at any plausible level of significance and, if anything, the point estimates are slightly larger for older users, who plausibly face more stable income processes.

5.2 Liquidity Constraints? Deposits over Income Flows

Next, we consider the potential role of liquidity constraints in explaining our results (Deaton, 1991). For this analysis, we split the sample by quintiles based on the deposits-to-income ratio in the month before activation. We use deposits-to-income as a proxy for liquidity constraints. If liquidity constraints explained our results, we would expect that users with lower deposits-to-income ratios would increase their spending and tap into negative deposits after accessing the overdraft facility, whereas otherwise similar users with high deposits-to-income ratios would barely react.

Figure 4 reports the results for estimating the interaction coefficients across quintiles of deposits-to-income. We detect substantial heterogeneity, but the heterogeneity of the effects goes in the *opposite* direction to what liquidity constraints would predict. In fact, the effect of overdraft availability on spending is (insignificantly) negative for users in the bottom quintile—the most liquidity-constrained users—and is zero for those in the second quintile. The effect increases nonlinearly and is disproportionally higher the higher the quintile. The estimated effect is about 2% for users in the third quintile, 5% in the fourth quintile, and 12% in the top quintile. It is the *least* liquidity-constrained users who increase their average spending more after having access to the overdraft facility relative to before.

This pattern is not an anomaly peculiar to users in the top quintile by liquidity: the positive effect is economically and statistically significant for the median user, and

is large also for users in the third and fourth quintiles, which suggests the effect we uncover involves the majority of the sample. We also want to stress that while liquidity-constrained users, the users in quintile 1, do not increase their consumption on average after the activation, they do so on impact in the month post activation and decrease their spending slightly in subsequent months (see Figure 5).

In Figure 5, we report the estimated coefficients across the deposit-to-inflows quintiles over time. The blue dashed line reports the estimated coefficients for users in the top quintile by deposits-to-inflow, the red solid line refers to users in the first quintile, and grey lines refer to users in other quintiles. Users in every quintile increase their consumption spending on impact to the availability of the credit line. However, we also see that users in the top quintile increase their spending more than others on impact and that these users increase their spending permanently, i.e. their spending does not revert to the pre-activation level over time. We do detect the same permanent effect for users in the fourth quintile, whereas the increase in spending reverts to zero for other users. Importantly, we detect no difference in spending patterns across users before the overdraft facility becomes available.

These consumption patterns suggest that the more liquid users might decrease their savings rate after accessing the overdraft facility, and that this decrease is permanent rather than transitory. We assess this mirroring conjecture by plotting the dynamics of the savings rate across quintiles of liquidity (see Figure IA1 in the Internet Appendix).[8] And, indeed, we find that the savings rate of the most liquid users drops permanently after they get access to the overdraft facility. We do not observe the same permanent pattern for the least liquid users, as well as for users in all other quintiles except for the fourth liquidity quintile.

5.2.1 Robustness

At first sight, the results in Figures 4, 5, and IA1 are puzzling: The association between users' pre-overdraft liquidity and their spending after overdraft activation seems to go in the *opposite* direction of what we would have expected under a liquidity-constraints explanation. Before digging deeper into the drivers of this empirical regularity, we perform a large set of robustness tests to verify this pattern is a robust feature of the data.

First, in Figure 6, we assess the robustness of the consumption pattern by liquidity quintile across a set of relevant subsamples and specifications. In the top left panel, we

[8] We thank Greg Buchak for suggesting that we directly compute the savings rate and examine its evolution over time.

compare the estimated coefficients and confidence intervals for the baseline specification and a second specification in which we additionally include interaction terms for (*i*) user age, (*ii*) inflow growth in the six months before and after overdraft availability, and (*iii*) inflow volatility in the twelve months prior to the overdraft treatment. We cannot reject the null that the coefficients are equal across specifications either economically or statistically.

In the top right panel of Figure 6, we compare the baseline effect in the full sample of users to the results when we restrict the working sample only to users with an active account for at least 12 months before the overdraft became available. This test is relevant because one might be concerned that users who drive our effect were merely opening an account to take advantage of the overdraft facility. These users might plan to make larger purchases and moved money from other accounts into the online account resulting in large savings-to-deposits and large spending responses. As expected, the estimates are substantially noisier in the smaller sample of long-term users, but we replicate our baseline results.

The middle panels of Figure 6 consider two forms of measurement error that would be problematic in our setting. First, we do not know whether the financial transactions we observe are all the transactions of our sample users because we cannot rule out for sure that these individuals have accounts at other banks. To alleviate this concern, we repeat the analysis restricting our sample to users whose average spending with our bank equals at least 40%, 50%, 60%, or 70% of their monthly income.[9] Our point estimates are similar and remain statistically significant when we increase this threshold despite the drop in the number of observations.[10]

Second, in the middle right panel of Figure 6, we consider the robustness of our results to the definition of our sorting variable—deposits-over-inflows. Specifically, we assess whether the consumption-liquidity pattern differs depending on whether we sort users based on the ratio measured in the month before the overdraft treatment or the one-month period starting 3 months before activation. We find no detectable difference in the baseline results which ensures transfers into the account because of future spending plans immediately before the activation cannot drive the patterns we document.

In the bottom panels, we show (*i*) that the statistical significance of our results barely changes across different levels of clustering and (*ii*) that our results remain robust when

[9] We thank Greg Buchak for suggesting this robustness test.

[10] As an additional (untabulated) test, we exclude all users who have any incoming or outgoing wire transfers other than direct income deposits, which might capture individuals who have other accounts at different institutions. Even in this case, the results barely change.

we restrict the sample to users who activate the overdraft at the first introduction of the overdraft product by the bank to alleviate the concern that users open accounts with the intention to active the overdraft.[11]

To further assess the robustness of the baseline consumption responses based on deposits-to-inflows, in Figure 7 we compare the estimated coefficients by quintiles across alternative demographic groups to check whether observables might drive the baseline relationship. This test is relevant because it could still be the case that individuals in the top quintiles by deposits-to-income might have characteristics that according to the PIH would predict a larger consumption response to the availability of credit despite the fact we do not observe that these characteristics matter unconditionally. Irrespective of whether we compare users who (*i*) live in European regions (NUTS 1) with above or below median savings rates (top left panel), (*ii*) reside in former Communist countries or not (top right panel), (*iii*) live in rural versus urban areas (bottom left panel) [12], or (*iv*) are young versus old (bottom right panel), we fail to detect any systematic differences in the baseline pattern.

Overall, the fact that users with higher deposits-to-inflow ratios spend more after they have access to the overdraft is a robust feature of the data.

5.3 Evidence on Objective Drivers of Precautionary Savings

In the standard life-cycle model, several individual characteristics, which are typically unavailable in observational field data, could predict the spending reaction to overdraft availability. For instance, systematic differences in users' risk aversion, time preferences, beliefs about life expectancy, or bequest motives might drive these results. Moreover, expecting imminent large expenses due to unobserved health reasons, other unobserved reasons, as well as private information about the probability of becoming unemployed and hence having more volatile income streams in the future might all contribute to the observed consumption effects (Gomes and Michaelides, 2005; Chetty and Szeidl, 2007; Carroll, 1997).

These dimensions either predict larger precautionary savings motives due to preference heterogeneity for a given level of objective uncertainty or might translate into differences in precautionary savings due to subjective needs. Once the overdraft facility becomes available to these users, it might act as a form of insurance and trigger higher spending,

[11] We thank Michaela Pagel for suggesting these robustness tests.

[12] We define "urban" users as those who live in cities with more than 500,000 inhabitants.

because users know they can now tap into negative deposits using the overdraft facility if a shock occurs.

In addition to these precautionary savings motives, one can think of other channels that could explain users' willingness to spend liquid assets after they obtain access to the overdraft. For instance, agents might have low financial literacy and might be unable to optimize their allocation of resources over time. Alternatively, users might be distrustful of others and hence display a precautionary savings motive that is not based on their financial characteristics.

In Figure 8, we consider two dimensions that might predict differential precautionary savings motives across liquidity quintiles, which we can directly measure in our data. First, we consider age in the top panel. Age increases monotonically with the bins by deposits-to-income inflows, ranging from 31.4 to 35.1. 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 age in the top bin. More 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 paths, might expect higher future income growth, might face less employment stability, and might still not be in the workforce at all.

Second, we test whether income uncertainty drives our results in the bottom panel. Specifically, we average the standard deviation of income inflows over the 12 months before overdraft facility activation within bins. We do not detect any systematic pattern or difference between the bottom and top bin by deposits-to-income. This fact is prima facie evidence that differences in income uncertainty do not justify why users in the top bin behave as if they have stronger precautionary savings motives.

We also report descriptive statistics by bins of deposits-to-inflows in Panel A of Table 3. We do not find large differences in demographics such as age or gender. Moreover, we find that lower outflows drive the higher deposits of users with high liquidity rather higher inflows or higher levels of deposits twelve months ago.

5.4 Eliciting Preferences, Beliefs, and Motivations

Considering the other dimensions we discuss above requires eliciting users' preferences, beliefs, and motivations. To make progress on this front, in June and July 2019, we fielded a survey intervention we designed to elicit users' risk and time preferences, a large set of beliefs, motivations, perceptions, as well as financial literacy and generalized trust.

Users could answer the elicitation survey either when using the app on their desktop or on their mobile phones. We sent 73,000 survey invitations, targeting a response rate between 10% to 15%, based on other surveys the FinTech bank ran on the platform in the recent past. Overall, we obtained 7,901 responses to the survey, which represents a response rate of 11%, in line with survey-based studies. Of the 7,901 respondents, we kept the survey outcomes of users for which we observe the activation of the overdraft facility during our sample period. Overall, we obtain 597 survey responses from users that belong to the baseline consumption sample and have activated the overdraft facility within our time frame.[13]

Importantly, when we approached users, we did not disclose that the aim of the intervention was to link their answers to the spending behavior around the introduction of the overdraft facility. Being silent about the overdraft facility was crucial to reduce the concern of experimenter demand effects—the possibility that subjects guess the aim of the experiment and align their answers and choices correspondingly.

The issue of experimenter demand effects is one of the strongest concerns faced by experimental economics and survey-based elicitation tools, as discussed recently, for instance, by De Quidt et al. (2018). Because we did not refer to the overdraft facility in any way within our survey, it is implausible that users would understand that the scope of the experiment was to assess the drivers of their spending and borrowing behaviors and hence could manipulate their answers to our questions. A drawback of the lack of direct reference to the overdraft facility experience is that we could not ask users directly for the motivations they had to activate and use the facility. Because of the issue of demand effects, the answers to such a direct question would have anyway been barely interpretable due to strategic motives in users' answers (see, e.g., D'Acunto, 2018, 2019).

We designed the survey questions following earlier research. We elicited risk aversion by asking users to rank their willingness to take on risks in financial matters on a scale from 1 (very low) to 10 (very high) (see, e.g, Guiso et al., 2008; D'Acunto et al., 2019b,c). For time preferences, we proposed users a hypothetical choice between a certain amount at the time of the survey or increasing amounts one month later (see, e.g., Benjamin et al., 2010; Coibion et al., 2019). To elicit expectations about upcoming large expenses, we asked users to rank their probability of facing large consumption expenses or large medical expenses over the following 12 months (D'Acunto et al., 2019a). We elicited similar rankings for users' expectations about the possibility of losing their job over the following 12 months—which aims to capture uncertainty about future income flows (Guiso

[13]We report the original survey questions in Internet Appendix Table IA2.

et al., 1992)—users' satisfaction with their health conditions, and users' generalized trust towards others (see, e.g., Dominitz and Manski, 2007; Guiso et al., 2004; D'Acunto et al., 2018). To elicit users' financial literacy, we asked them to assess whether the amount that would compound in their checking account at a certain interest rate would be above or below a given value (see, e.g., Lusardi and Mitchell, 2011). Among the potential questions the literature proposed to assess financial literacy (Lusardi and Mitchell, 2014), the ability to understand compounding is most relevant for our setting, in which we study the borrowing and spending behavior of bank customers that consider the use of a credit line.

Panel B of Table 3 reports the results for the basic preferences, beliefs, and motivations we elicited. We sort users into five quintiles based on their deposit-to-income ratio before accessing the overdraft facility. We then compute the average quantitative response of these users to the elicitation questions within each bin, which we plot together with the 95% confidence interval of the mean in each quintile in Figure 9.

Across the dimensions we elicited, we fail to detect any systematic patterns in the cross-section of liquidity. None of the dimensions we consider, which could have explained the responses in Figure 4 in a standard life-cycle consumption–savings model, seems able to capture such patterns. Not only are the averages within quintiles not different from each other statistically, but they are also similar in terms of economic magnitude, which suggests that low statistical power in the small survey sample is unlikely to drive the lack of variation across quintiles.

Overall, our evidence based on elicited preferences, beliefs, and motivations directly dismisses the most compelling standard channels that would justify a spending increase by liquid users, but not illiquid users, once the overdraft facility becomes available to them.

5.5 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 spending. This pattern is intriguing, because we might have expected the most liquid users to be those that had the least need of an overdraft facility if they wanted to spend more before activation. In contrast, if the overdraft facility is mainly used to smooth spending and alleviate liquidity constraints, users in the bottom liquidity quintile should react most.

Alternatively, users might accumulate liquidity over time because of the perception of possible future needs. Therefore, sorting users at the time of the overdraft activation corresponds to sorting users based on their perceived need of precautionary savings. So far, comparing bins by deposits-to-income ratios does not suggest that users in the top bin have any objective reason to have stronger precautionary savings motives than users in the lower bins. We now 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 using the overdraft. Precautionary savers, contrary to liquidity-constrained individuals, would likely not tap into negative deposits and would not increase their debt levels through the overdraft facility. Instead, they should view the facility as a form of insurance against negative income shocks or unexpected expenses and would thus spend some of the existing liquidity they had accumulated before the overdraft facility was available once they know they can tap into negative deposits if needed. Hence, we would expect a permanent increase in the consumption spending of these high deposit-to-income users and a decrease in the savings rate as we indeed observe in the data.

The results in Figure 4 are broadly consistent with users in the top liquidity bin 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 these other users. The probability of tapping into negative deposits in the three months after activation ranges from 67% for users in the bottom bin to 10% for users in the top bin. Moreover, it is unlikely that a large fraction of the individuals in our sample have objective precautionary savings motives due to potentially unexpected large medical bills or other medical-related expenses, because the 75th percentile of age in our sample equals 38.7 years.

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. However, in the data we do not observe heterogeneity in the cumulative inflows at our bank in the three months before activation by deposits-to-income. Moreover, we test directly for this potential alternative interpretation by comparing the change in the ratio of spending on big ticket items over total spending before and after activation of the overdraft facility and across users by liquidity quintiles.[14] Figure IA2

[14]For this test, we define as big ticket items purchases of goods and services which cost above 5,000 Euros. The results are similar if we change this threshold.

in the Internet Appendix reports the change in this ratio across liquidity quintiles, and unveils that, in fact, the ratio has barely changed economically and statistically for most users. In particular, the change does not display the monotonic pattern by liquidity that we should observe if this explanation was relevant to our results.

Overall, users with a high share of deposits to income and hence high liquidity behave as if they had precautionary savings motives and hence accumulated savings and saved a larger share of their income 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 expenses 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 or higher risk aversion. Based on these considerations, we label the mechanism we document in this paper as *perceived* precautionary savings.

5.6 Alternative Explanations and Channels

We move on to consider a set of additional potential alternative interpretations of our results. First, the facility might free up liquid resources users were keeping in their bank accounts because of private information about upcoming individual income shocks. In this vein, even models of buffer-stock savings that allow for both impatience and precautionary savings motives might explain our results at least in part. 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 the reaction to the availability of credit. Also, as discussed above, users close to the liquidity constraint do not react on average post activation, whereas those farther away from the constraint react the most, and the buffer-stock interpretation predicts the opposite pattern.

A second potential explanation is a buffer-stock model with durable consumption similar to the one Aydin (2015) studies. However, a set of results suggests this interpretation cannot fully explain the results in our setting. In addition to the facts that the least liquidity-constrained users react the 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 third interpretation we consider is present-biased preferences—the fact individuals

discount the distant future by more than they discount the immediate future (Meier and Sprenger, 2010). 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 to the activation of the overdraft facility.

At first, our results might appear inconsistent with a large literature documenting users that are ex-ante most constrained react the most to the extension of credit (see, e.g., Agarwal et al., 2017; Aydin, 2015; Gross and Souleles, 2002). The major difference between these studies and our paper is that earlier studies exploited changes in the intensive margin of credit, that is, changing the extent of credit for agents who were already borrowing. We, instead, focus on the extensive margin of credit, that is, providing users with a borrowing facility for the first time. The insurance effect of first-time borrowing facilities could not be detected in earlier research even if this effect was true, because all the agents in earlier work were already provided with insurance around the change in credit limits. The insurance mechanism did not change within agents over time in intensive-margin studies reconciling our work with previous findings in the literature.

In our setting, users that previously were most concerned about future unexpected expenses or had higher precautionary savings demands for other reasons that we capture by sorting on deposits-to-income seem to react the most to the provision of a downside insurance mechanism—the overdraft facility.

6 Regression Discontinuity Analysis

An important identification concern with our analysis so far is the endogenous timing at which users open their accounts. One might worry that users activate the account only when they foresee the need for an overdraft facility, for instance when they plan to have larger-than-usual upcoming expenses.

We provide tentative evidence that this endogeneity concern might not be as relevant in our setting as one might worry. To do so, we repeat our tests by restricting the sample only to users who activated their overdraft facility within one month after the online

bank made it available.[15] The rationale for this test is that users who activate the facility as soon as it is available might be less likely to have planned to activate it in view of large upcoming expenses, because they could not know the fact that the bank would have provided such facility, nor the timing. We report the results for this test in Figure 6 (bottom left panel). Although this test restricts the baseline sample substantially, and hence increases our estimated standard errors, the baseline spending patterns by liquidity quintiles are quite similar to the baseline analysis in terms of economic magnitudes, and the coefficients for the two top quintiles, as in the baseline analysis, are statistically different from zero.

Despite this evidence, one might worry that other 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, even for users who have been banking with our provider for quite some time and hence do not open new accounts.

To directly address these endogeneity concerns, in this last part of the paper we estimate the causal effects of the availability of the mobile overdraft on spending in a sharp RDD that exploits variation in users' overdraft limits based on thresholds embedded in the bank's credit allocation algorithm. Our sharp RDD conditions the analysis on users' (possibly endogenous) selection into the mobile overdraft and relies on exogenous variation in the size of the credit line along the intensive margin.

We do not implement this design for the baseline analysis of the paper because we can only exploit a limited number of users in the sample for the RDD setting. By construction, as we discuss in more detail below, the RDD only uses users who are close to pre-specified thresholds based on the algorithm that assigns the overdraft amount to those activating an overdraft facility. Because of the limited sample, any meaningful statistical analysis of heterogeneity and variation across subgroups of users would be impossible.

6.1 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

[15] We thank Michaela Pagel and Greg Buchak for suggesting this test.

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:

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

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., pensions, 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 10 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. Crucially, users are not aware of the algorithm and the algorithm is fully automated leaving no room for human intervention.

6.2 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 RDD 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 10 illustrates the exact treatment rule of our sharp RDD 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 RDD:

$$\tau \equiv \mathbb{E}\left(C_i(1)|X_i = 0\right) - \mathbb{E}\left(C_i(0)|X_i = 0\right). \tag{3}$$

τ is the RD treatment effect and $C_i(1/0)$ is the change in treated (1) or control (0) user's average consumption three months before and after the credit allocation decision, divided by the individual's average inflows in the three 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 two local weighted regressions. Formally, each user's consumption change

equals:

$$C_i = \begin{cases} C_i(0) & \text{if } X_i \geq 0 \\ C_i(1) & \text{if } X_i < 0. \end{cases} \tag{4}$$

We focus on observations within the interval $[-h, h]$ around the rounding threshold, where $h > 0$ denotes our bandwidth. 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:

$$\hat{C}_i = \hat{\mu}_{+/-} + \sum_{j=1}^{p} \hat{\mu}_{+/-,j} X_i^j. \tag{5}$$

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

$$\hat{\tau} = \hat{\mu}_+ - \hat{\mu}_-. \tag{6}$$

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 bias resulting from subjective bandwidth choices. We residualize the outcome variables of our RD analysis with NUTS3×year-month fixed effects to ensure that we compare treated and control users from the same European country at a similar point in time.

6.3 Assessing Identification Assumptions: Treatment Manipulation and Balancing Tests

Our sharp RDD 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 users can control their treatment assignment in our set-

ting. Most importantly, the bank's credit allocation algorithm is proprietary information and overdraft users do not know it. Even if individuals were informed about the precise inner workings of the overdraft allocation formula (in particular its rounding thresholds), it seems implausible they could precisely manipulate their income, for example, by negotiating a higher wage with their employer (Lee and Lemieux, 2010).

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 to the right of 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 Appendix Figure IA3, 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 Figure IA4 in the 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 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, no systematic differences in observable characteristics between the two groups of individuals should exist. Consistent with this argument, we do not find significant differences in the average age, gender, time since account opening, 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 IA1 of the Internet Appendix, we document 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 the local continuity assumption is satisfied and thereby corroborates the internal validity of our sharp RD design.

6.4 RD Consumption Effect

In Figure 11, 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 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 two linear regressions through all observations on each side of the rounding threshold.

In Panel A of Figure 11, 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 after and before the user obtained access to the credit line. We find a positive discontinuity in users' consumption growth right at the rounding threshold, indicating treated users consume more relative to control users following the exogenous assignment of a 250 Euro higher overdraft limit. In Table 5, we present the coefficients from estimating the sharp RD design we formalized in equations (3) to (6). 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. 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. (2019), we find adding covariates increases the precision of our point estimates, which again suggests the local continuity assumption is satisfied.

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 IA5 in the Internet Appendix shows 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 (see, 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). Internet Appendix Figure IA6 plots the coefficient estimates for different choices of r and shows that observations close to the rounding threshold do not drive our results. Taken together, the RD results confirm our baseline findings and alleviate concerns that selection into the app to obtain access to a line of credit in the anticipation of upcoming expenses drive our results.

7 Conclusion

We study the consumption response to the introduction of an overdraft facility on a FinTech app. The average user increases her consumption spending over income by 4.5 percentage points on impact relative to similar users that have access to the overdraft facility at a later point in time. The increase in consumption has a permanent component and we observe a reallocation of consumption from discretionary to non-discretionary expenses.

The surprising result is that we find a large spending response for users with high liquid savings, whereas the users with the lowest amount of liquid savings barely react at all to the provision of the overdraft facility. We do not detect any heterogeneity in preferences or beliefs across liquidity quintiles, nor heterogeneity in standard theoretical drivers of precautionary savings motives. These results are not fully consistent with myopic consumers, models with financial constraints, buffer stock models (with durables) and the canonical life-cycle permanent income model.

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 steep and flat income paths.

We argue that a *perceived* precautionary savings channel might be 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 in fact barely make use the overdraft facility.

Our findings open new paths 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, can modified versions of the neoclassical consumption model predict the patterns we uncover? 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

Agarwal, S., Chomsisengphet, S., Mahoney, N., Stroebel, J., 2017. Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow? Quarterly Journal of Economics 133, 129–190.

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

Bain, 2017. Evolving the Customer Experience in Banking. https://www.bain.com/insights/evolving-the-customer-experience-in-banking, accessed on january 3, 2019. Bain & Company Inc.

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

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

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

Ben-David, I., Fermand, E., Kuhnen, C.M., Li, G., 2018. Expectations Uncertainty and Household Economic Behavior. Working Paper. National Bureau of Economic Research.

Benjamin, D.J., Choi, J.J., Strickland, A.J., 2010. Social Identity and Preferences. American Economic Review 100, 1913–28.

Calonico, S., Cattaneo, M.D., Farrell, M.H., Titiunik, R., 2019. Regression Discontinuity Designs Using Covariates. Review of Economics and Statistics 101, 442–451.

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

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

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

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

Cheng, M.Y., Fan, J., Marron, J.S., 1997. On Automatic Boundary Corrections. Annals of Statistics 25, 1691–1708.

Chetty, R., Szeidl, A., 2007. Consumption Commitments and Risk Preferences. Quarterly Journal of Economics 122, 831–877.

Coibion, O., Gorodnichenko, Y., Weber, M., 2019. Monetary Policy Communications and their Effects on Household Inflation Expectations. Working Paper. National Bureau of Economic Research.

D'Acunto, F., 2018. Identity and Choice Under Risk. Working Paper. Boston College.

D'Acunto, F., 2019. Tear Down This Wall Street: Anti-finance Rhetoric, Subjective Beliefs, and Investment. Working Paper. Boston College.

D'Acunto, F., Hoang, D., Paloviita, M., Weber, M., 2019a. IQ, Expectations, and Choice. Working Paper. National Bureau of Economic Research.

D'Acunto, F., Malmendier, U., Ospina, J., Weber, M., 2019b. Exposure to Daily Price Changes and Inflation Expectations. Working Paper. National Bureau of Economic Research.

D'Acunto, F., Malmendier, U., Weber, M., 2019c. Traditional Gender Norms Distort Women's Economic Expectations. Working Paper. Booth School of Business.

D'Acunto, F., Prabhala, N., Rossi, A.G., 2019a. The promises and pitfalls of robo-advising. The Review of Financial Studies 32, 1983–2020.

D'Acunto, F., Prokopczuk, M., Weber, M., 2018. Historical Antisemitism, Ethnic Specialization, and Financial Development. Review of Economic Studies 86, 1170–1206.

D'Acunto, F., Rossi, A.G., 2020. Robo-advising. Palgrave Macmillan Handbook of Technological Finance .

D'Acunto, F., Rossi, A.G., Weber, M., 2019b. Crowdsourcing financial information to change spending behavior .

De Quidt, J., Haushofer, J., Roth, C., 2018. Measuring and Bounding Experimenter Demand. American Economic Review 108, 3266–3302.

Deaton, A., 1991. Saving and Liquidity Constraints. Econometrica 59, 1221–48.

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

Dominitz, J., Manski, C.F., 2007. Expected Equity Returns and Portfolio Choice: Evidence from the Health and Retirement Study. Journal of the European Economic Association 5, 369–379.

Friedman, M., 1957. The Permanent Income Hypothesis, in: A Theory of the Consumption Function. Princeton University Press, pp. 20–37.

Gargano, A., Rossi, A.G., 2018. Does it pay to pay attention? The Review of Financial Studies 31, 4595–4649.

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

Gomes, F., Michaelides, A., 2005. Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence. Journal of Finance 60, 869–904.

Gross, D.B., Souleles, N.S., 2002. Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data. Quarterly Journal of Economics 117, 149–185.

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

Guiso, L., Jappelli, T., Terlizzese, D., 1992. Earnings Uncertainty and Precautionary Saving. Journal of Monetary Economics 30, 307–337.

Guiso, L., Sapienza, P., Zingales, L., 2004. The Role of Social Capital in Financial Development. American Economic Review 94, 526–556.

Guiso, L., Sapienza, P., Zingales, L., 2008. Trusting the Stock Market. Journal of Finance 63, 2557–2600.

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

ING, 2018. International Survey Mobile Banking 2018. https://www.ezonomics.com/ing_international_surveys/mobile-banking-2018, accessed on january 3, 2019. ING-DiBa AG.

Jappelli, T., Pistaferri, L., 2017. The Economics of Consumption: Theory and Evidence. Oxford University Press.

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

Lee, D.S., Lemieux, T., 2010. Regression Discontinuity Designs in Economics. Journal of Economic Literature 48, 281–355.

Lusardi, A., Mitchell, O.S., 2011. Financial Literacy and Planning: Implications for Retirement Wellbeing. Working Paper. National Bureau of Economic Research.

Lusardi, A., Mitchell, O.S., 2014. The Economic Importance of Financial Literacy: Theory and Evidence. Journal of Economic Literature 52, 5–44.

Meier, S., Sprenger, C., 2010. Present-Biased Preferences and Credit Card Borrowing. American Economic Journal: Applied Economics 2, 193–210.

Mian, A., Sufi, A., 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., Sufi, A., Verner, E., 2017. Household Debt and Business Cycles Worldwide. Quarterly Journal of Economics 132, 1755–1817.

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

Roberts, M.R., Whited, T.M., 2013. Endogeneity in Empirical Corporate Finance, in: Handbook of the Economics of Finance. Elsevier. volume 2, pp. 493–572.

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

Tables and Figures

Figure 2: Consumption Pattern around Overdraft Availability

Notes: 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: Intertemporal Consumption Smoothing?

Notes: 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.

Figure 4: Consumption Response by Deposit-to-Income Quintile

Notes: This figure illustrates the cross-sectional heterogeneity in users' consumption response to the mobile overdraft. To generate this plot, 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.

Figure 5: Consumption Pattern by Deposit-to-Income Quintile

Notes: This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the heterogeneous effect of mobile overdrafts on the consumption behavior of users with different ex-ante liquidity. To generate this plot, 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 deposit to inflows ratio in the month before treatment. We estimate model (1) from Table 2 but replace the *Overdraft Available* indicator with separate time dummies, which we further interact with our quintile indicators. Each time dummy marks a one-month period (except for event period $t-1$). Coefficients are normalized by subtracting the pre-treatment mean for each quintile. We double cluster standard errors at the NUTS2 and year-month level.

Figure 6: Robustness Tests

Notes: This figure illustrates the cross-sectional heterogeneity in users' consumption response to the mobile overdraft for various robustness specifications. 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 deposit to inflows ratio. 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. In the top left panel, we additionally control for user characteristics at treatment (overdraft available interacted with user age, inflows growth and inflows volatility at treatment). In the top right panel, we keep only users that had an account with the bank for at least 12 months before they activated the overdraft facility. In the middle left panel, we focus on users who exhibit an average consumption rate of at least $X\%$ over the sample period where we vary X from 40 to 70. In the middle right panel, we group users into quintiles based on the average deposit to income ratio from 12 until 3 months before treatment. In the bottom left panel, we focus on users who activate the overdraft at the introduction date and compare them to users who activate the overdraft at least one year after the introduction. In the bottom right panel, we provide baseline regression results for various standard error clusterings. Vertical bands represent 95% confidence intervals for the point estimates of each quintile. Except for the bottom right panel, we double cluster standard errors at the NUTS2 and year-month level.

Figure 7: External Validity

Notes: This figure illustrates the cross-sectional heterogeneity in users' consumption response to the mobile overdraft for various sample splits. 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 (quintile) based on the deposit to inflows ratio. 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. Note: savings rate refers to the NUTS1 average savings rate in 2015 from statistisches Bundesamt.

Figure 8: User Characteristics by Deposit-to-Income Quintile

Notes: This figure illustrates the cross-sectional heterogeneity in users' characteristics by deposits to inflows quintile. 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 deposit to inflows ratio in the month before treatment. Vertical bands represent 95% confidence intervals for the mean of each quintile.

Figure 9: Subjective Beliefs by Deposit-to-Income Quintile

Notes: This figure plots a set of preferences and beliefs dimensions we elicited from users through an ad-hoc survey intervention. 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 ratio of deposited amount over income. We limit the sample to users that activate the overdraft facility in our main sample. Vertical bands represent 95% confidence intervals for the point estimates of each quintile.

Figure 10: Treatment Assignment in Regression Discontinuity Analysis

Notes: 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.

Figure 11: Regression Discontinuity Results

Panel A: Overdraft Amounts

Panel B: Consumption Growth

Notes: 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 12 disjoint bins, calculate the average value, plot this value above the midpoint of each bin, and separately fit two linear regressions through all observations on each side of the rounding threshold.

Table 1: Descriptive Statistics

	N	Mean	SD	P10	P25	P50	P75	P90
User Age [Years]	603,157	34.07	9.93	23.91	27.00	31.49	38.66	48.91
Female [0/1=Yes]	603,157	0.21	0.41	0.00	0.00	0.00	0.00	1.00
Urban [0/1=Yes]	603,157	0.53	0.50	0.00	0.00	1.00	1.00	1.00
Account Age [Years]	603,157	1.57	0.86	0.47	0.88	1.49	2.18	2.78
Overdraft Available [0/1=Yes]	603,157	0.91	0.29	1.00	1.00	1.00	1.00	1.00
Overdraft Amount [Euro]	574,328	1,561.38	1,602.05	500.00	500.00	1,000.00	2,000.00	3,000.00
Consumption [Euro]	603,157	730.81	949.42	0.00	131.00	465.00	993.00	1,715.00
Inflows [Euro]	603,157	2,219.84	5,948.55	90.00	380.00	1,124.00	2,534.00	4,564.00
Consumption / Inflows [%]	603,157	48.11	44.89	0.00	13.88	37.24	70.08	108.00
Card Consumption [Euro]	603,157	498.58	736.64	0.00	67.00	281.00	646.00	1,206.00
Card Consumption / Inflows [%]	603,157	33.80	37.17	0.00	6.49	22.12	47.85	84.72
Cash Withdrawals [Euro]	603,157	220.92	400.55	0.00	0.00	90.00	300.00	600.00
Cash Withdrawals / Inflows [%]	603,157	13.44	19.56	0.00	0.00	5.33	19.27	40.00
Discretionary [Euro]	603,157	208.36	419.58	0.00	1.00	70.00	242.00	550.00
Non-Discretionary [Euro]	603,157	290.22	469.75	0.00	26.00	159.00	373.00	692.00
Discretionary / Non-Discretionary [%]	484,683	61.03	79.26	0.00	7.66	34.41	82.45	155.17

Notes: This table reports descriptive statistics for our user-month panel. The sample consists of 36,005 users that received an overdraft between February 2015 and September 2017 and covers each individual's complete transaction history from February 2015 to March 2019. 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 A1.

Table 2: Effect of Overdraft Availability on Users' Consumption Behavior

Dependent Variable (×100):	$\frac{\text{Consumption}_t}{\text{Inflows}_{t-1}}$	$\frac{\text{Card Consumption}_t}{\text{Inflows}_{t-1}}$	$\frac{\text{Cash Withdrawals}_t}{\text{Inflows}_{t-1}}$	$\frac{\text{Discretionary}_t}{\text{Non-Discretionary}_t}$
	(1)	(2)	(3)	(4)
Overdraft Available$_t$	4.572***	3.152***	1.345***	2.902***
	(11.11)	(9.71)	(7.73)	(5.38)
Economic Effect Size:	9.503%	9.325%	10.007%	4.755%
Fixed Effects:				
User	Yes	Yes	Yes	Yes
NUTS3 × Year-Month	Yes	Yes	Yes	Yes
Standard Error Clusters:				
NUTS2	48	48	48	48
Year-Month	49	49	49	50
Adjusted R^2	0.301	0.300	0.339	0.185
User-Year-Month Observations	603,157	603,157	603,157	484,683

Notes: 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 withdrawals from ATMs in the given month. *Inflows* is the total amount of all incoming transactions a user receives in the given month. *Discretionary* is the sum of users' monthly spending on Entertainment, Shopping, Gastronomy, and Travel. *Non-Discretionary* consumption equals Card Consumption minus Discretionary spending. *Overdraft Available* is a binary indicator that equals 1 if the user has access to a mobile overdraft in the given month. We compute the *Economic Effect Size* as the pre-treatment average of each outcome variable multiplied by the corresponding coefficient estimate. 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.

Table 3: User Characteristics by Deposit-to-Income Quintile

Panel A: Users' Financial Characteristics

	Deposits / Inflows Quintiles					
	Q1	Q2	Q3	Q4	Q5	SD
Flow Characteristics						
Cumulative $\text{Inflows}_{t-12:t-1}$	11,491.77	13,381.19	14,930.03	10,580.58	10,753.34	16,980.36
Inflows_{t-1}	1,270.25	1,775.88	1,882.12	1,779.34	1,293.43	3,961.29
$\text{Inflows}_{t-12:t-1}$	957.65	1,115.10	1,244.17	881.71	896.11	1,415.03
$\text{SD}(\text{Inflows}_{t-12:t-1})$	836.63	941.74	1,191.37	793.55	901.77	1,635.70
$\text{SD}(\text{Consumption}_{t-12:t-1})$	462.59	471.12	483.49	473.75	429.53	476.92
Cumulative $\text{Outflows}_{t-12:t-1}$	11,778.41	13,135.22	13,767.61	9,589.49	7,774.88	15,500.42
Deposit Characteristics						
Deposits_{t-1}	39.45	451.87	1,128.83	1,715.09	4,163.05	8,520.52
Deposits_{t-12}	283.02	320.74	468.30	434.86	627.48	1,225.21
$\text{Deposits}_{t-12:t-1}$	273.90	459.46	729.29	814.87	2,012.25	2,432.76
Negative $\text{Deposits}_{t+1:t+3}$	0.65	0.41	0.29	0.19	0.10	0.47
Account Characteristics						
User Age_t	31.35	32.25	32.69	33.15	35.05	9.89
Female	0.20	0.20	0.22	0.25	0.23	0.41
Urban	0.50	0.54	0.56	0.54	0.47	0.50
Account Age_t	0.67	0.62	0.61	0.55	0.71	0.36
Users	1,909	1,909	1,910	1,908	1,909	9,545

Panel B: Users' Subjective Beliefs

	Deposits / Inflows Quintiles					
	Q1	Q2	Q3	Q4	Q5	SD
Risk Aversion [0-10]	5.13	4.63	4.47	4.65	4.72	2.31
Trust Level [0-10]	5.29	5.74	5.53	5.31	5.47	2.15
Risk-Return Category [1-4]	2.88	2.85	3.03	2.93	2.86	0.74
Patience [> 100]	118.98	119.07	119.45	121.96	117.48	17.42
Likelihood of Large Expenses [0-10]	4.13	4.29	4.48	4.46	4.80	2.63
Likelihood of Large Medical Expenses [0-10]	2.32	2.26	2.32	2.50	2.45	2.29
Likelihood of Loosing Job [0-10]	1.34	1.15	1.41	1.14	1.29	2.12
Health Satisfaction [0-10]	6.99	7.28	7.16	7.02	7.41	2.21
Reasons for Saving [1-4]	2.20	2.26	2.23	2.35	2.29	0.98
Basic Financial Computation Ability [1-3]	2.85	2.72	2.82	2.70	2.69	0.56
Life Expectancy [number]	83.37	85.57	84.70	83.06	83.13	10.63
Users	96	141	112	114	136	599

Notes: This table provides descriptive statistics of user characteristics and survey responses. Panel A provides descriptive statistics for users at the time of overdraft activation, sorted into quintiles based on the ratio of deposits over inflows in the month before activation of the overdraft facility. Subscript t refers to the activation date, while T denotes the last month of each user in our sample. The last column reports the standard deviation of the corresponding characteristic across all users. Panel B reports a set of preferences and beliefs dimensions we elicited from users through an ad-hoc survey intervention. 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 ratio of deposited amount over income. We limit the sample to users that activate the overdraft facility in our main sample.

Table 4: Assessing RDD Identification Assumptions

Panel A: Treatment Manipulation Tests

		Bandwidths		Effective N		Test	
		Left	Right	Left	Right	T-Stat.	p-value
$h_- \neq h_+$	$T_2(\hat{h}_1)$	104.17	76.21	373	328	-0.94	0.35
	$T_3(\hat{h}_2)$	75.73	82.98	285	345	-0.85	0.39
	$T_4(\hat{h}_3)$	89.01	55.32	322	221	0.39	0.70
$h_- = h_+$	$T_2(\hat{h}_1)$	111.32	111.32	424	448	-0.05	0.96
	$T_3(\hat{h}_2)$	71.59	71.59	248	289	-0.51	0.61
	$T_4(\hat{h}_3)$	58.32	58.32	201	229	-0.59	0.56

Panel B: User Characteristics around Rounding Thresholds

	Rounded Up			Rounded Down			Difference in Means	
	Mean	SD	Median	Mean	SD	Median	Diff.	t-Stat.
User Age [Years]	32.32	9.34	29.91	33.02	10.12	30.37	0.70	1.13
Female [0/1=Yes]	0.25	0.43	0.00	0.24	0.43	0.00	-0.01	-0.35
Account Age [Years]	0.87	0.38	0.83	0.85	0.41	0.79	-0.01	-0.53
Inflows$_{t-3:t-1}$ [Euro]	1,405.64	1,530.10	1,101.67	1,458.92	1,454.29	1,069.67	53.28	0.56
Consumption$_{t-3:t-1}$ [Euro]	558.89	522.39	427.17	620.04	597.38	482.83	61.14*	1.70
Users	500			474			974	

Notes: This table provides the results of tests assessing our RDD identification assumptions. Panel A 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 on each side of the threshold, while the tests in the last three rows impose a common bandwidth on both sides of the threshold. Panel B 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.

Table 5: Consumption Growth around Rounding Thresholds

Dependent Variable ($\times 100$):	$\frac{\text{Consumption}_{t+1:t+3} - \text{Consumption}_{t-3:t-1}}{\text{Inflows}_{t-3:t-1}}$			
	(1)	(2)	(3)	(4)
Conventional	17.63**	23.48**	17.96**	23.70**
	(2.19)	(2.20)	(2.12)	(2.34)
Robust	21.56**	26.52**	21.83**	26.62**
	(2.37)	(2.31)	(2.31)	(2.46)
Covariates	No	No	Yes	Yes
User Observations	876	876	876	876
Order Local Polynomial (p)	1	2	1	2
Order Bias (q)	2	3	2	3
Bandwidth Left	25.47	35.89	23.98	36.86
Bandwidth Right	25.47	35.89	23.98	36.86

Notes: 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, 2019). In columns (1) and (2), we do not add any covariates. In columns (3) and (4), we control for *User Age*, gender (*Female*), and *Account Age*. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Appendix

Table A1: Variable Definitions

Variable	Definition
Account Age	End of month date minus the date when the user completed the account opening procedure.
Big Ticket Expenditure	Indicator variable equal to one if the user has at least one transaction of more than 5,000 EUR in the given month.
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*.
Deposits	Total amount of available funds in a user's account at the end of a month.
Discretionary	Sum of users' monthly expenditures on *Entertainment*, *Shopping*, *Gastronomy*, and *Travel*.
Female	Indicator variable equal to one if the user is female.
Incoming Wire Transfers	Total amount of all SEPA credit transfers a user receives in the given month.
Inflows	Total amount of all incoming transactions a user receives in the given month.
Negative Deposits	Indicator variable equal to one if the user has a negative account balance in the given month.
Non-Discretionary	*Card Consumption* minus *Discretionary* spending.
Overdraft Available	Indicator variable equal to one if the user has access to a mobile overdraft in the given month.
Overdraft Amount	Maximum overdraft amount granted to the user by the bank in the given month.
Savings Rate	Lagged inflows minus total current outflows divided by lagged inflows.
Urban	Indicator variable equal to one if the user lives in a NUTS3 region with a population of at least 500,000 people.
User Age	End of month date minus first day of user's birth year.

Internet Appendix to

Perceived Precautionary Savings Motives: Evidence from FinTech

Francesco D'Acunto Thomas Rauter

Christoph Scheuch Michael Weber

Figure IA1: Savings Rate Pattern by Deposit-to-Income Quintile

Notes: This figure shows coefficient estimates and 95% confidence intervals for OLS regressions estimating the heterogeneous effect of mobile overdrafts on the savings behavior of users with different ex-ante liquidity. To generate this plot, 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 deposit to inflows ratio in the month before treatment. We estimate model (1) from Table 2 but replace the *Overdraft Available* indicator with separate time dummies, which we further interact with our quintile indicators. Each time dummy marks a one-month period (except for event period $t-1$). Coefficients are normalized by subtracting the pre-treatment mean for each quintile. We double cluster standard errors at the NUTS2 and year-month level.

Figure IA2: Big Ticket Expenditure Response by Deposit-to-Income Quintile

Notes: This figure illustrates the cross-sectional heterogeneity in users' big-ticket expenditure response after accessing the mobile overdraft. To generate this plot, 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.

Figure IA3: Distribution of Forcing Variable around Rounding Thresholds

Panel A: Histogram of Forcing Variable

Panel B: Local Polynomial Density Estimate

Notes: This figure provides graphical evidence for our treatment manipulation tests in Section 6.3. 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 forcing variable X_i around the rounding threshold.

Figure IA4: Robustness Analyses for Treatment Manipulation Tests

Notes: 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 IA5: Sensitivity of RDD Effects to Bandwidth Choice

Notes: This figure shows convetional regression (RD) discontinuity estimates (solid line) and corresponding 95% confidence intervals (shaded area) for varying bandwidth choices. The figure demonstrates 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 IA6: Sensitivity of RDD Effects to Observations around Rounding Thresholds

Notes: This figure shows robust regression discontinuity (RD) estimates (solid line) and corresponding 95% confidence intervals (shaded area) for varying donut hole radius choices. This figure demonstrates 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 IA1: User Characteristics around Rounding Thresholds

Dependent Variable:	User Age	Female	Account Age
	(1)	(2)	(3)
Conventional	3.485	-0.00263	0.0568
	(1.28)	(-0.02)	(0.61)
Robust	4.192	-0.0407	0.0432
	(1.26)	(-0.24)	(0.39)
Covariates	No	No	No
User Observations	972	972	972
Order Local Polynomial (p)	1	1	1
Order Bias (q)	2	2	2
Bandwidth Left	29.77	33.95	41.64
Bandwidth Right	29.77	33.95	41.64

Notes: 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. 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, 2019). ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Table IA2: Survey Questions

Question	Possible Answers
Do you rate yourself as a risk-taking person or are you trying to avoid risks?	Scale from 0 (not willing to take risks) to 10 (very willing to take risks)
Would you say that you are a person who trusts others, or not?	Scale from 0 (I do not trust others) to 10 (I trust others fully)
When you make savings or investment decisions for yourself, which of the following statements describes you best?	1 – I take significant risks and want to generate high returns. 2 – I take above-average risks and want to achieve above-average returns. 3 – I take average risks and aim for average returns. 4 – I am not willing to take any financial risk and accept not to generate any returns.
Imagine, you get either 100 Euros immediately or a higher amount of money in a month. What is the lowest amount you would be willing to wait for a month?	1 – 101 Euros 2 – 103 Euros 3 – 108 Euros 4 – 117 Euros 5 – 125 Euros 6 – 133 Euros 7 – 150 Euros 8 – 167 Euros 9 – 183 Euros 10 – 200 Euros 11 – 233 Euros
How likely is it that you will face large, unexpected expenses over the next 12 months?	Scale from 0 (very unlikely) to 10 (very likely)
How likely is it that you will face large medical expenses for yourself or a family member over the next 12 months, including hospitalization and nursing care?	Scale from 0 (very unlikely) to 10 (very likely)
How likely are you to lose your job in the next 12 months?	Scale from 0 (very unlikely) to 10 (very likely)
How satisfied are you with your health at the moment?	Scale from 0 (not satisfied) to 10 (very satisfied)
What do you usually think of when saving?	1 – I usually save for very specific expenses, such as a vacation or a car. 2 – I usually save to have a small amount of money available for unexpected expenses. 3 – I do not save much or cannot save much. Whenever I put money aside, I do it without much thought. 4 – I am saving with the goal of building a small estate that I may pass on to my children, nephews, nieces, or other family members.
Suppose you have 100 Euros in your savings account and earn 10 percent interest per year. How high will be your balance after 2 years?	1 – Lower than 120 Euros 2 – 120 Euros 3 – Higher than 120 Euros
How high do you estimate your life expectancy?	Enter number

Table IA3: Characteristics of Survey Respondents

	Non-Respondents	Respondents	
	Mean	Mean	SD
User Age [Years]	33.78	38.37	9.93
Female [0/1=Yes]	0.22	0.12	0.41
Urban [0/1=Yes]	0.54	0.39	0.50
Account Age [Years]	1.56	1.77	0.86
Overdraft Available [0/1=Yes]	0.91	0.91	0.29
Overdraft Amount [Euro]	1,566.51	1,484.89	1,602.05
Consumption [Euro]	742.09	561.75	949.42
Inflows [Euro]	2,236.78	1,965.94	5,948.55
Consumption / Inflows [%]	48.33	44.86	44.89
Card Consumption [Euro]	506.49	379.95	736.64
Card Consumption / Inflows [%]	33.91	32.22	37.17
Cash Withdrawals [Euro]	223.92	175.96	400.55
Cash Withdrawals / Inflows [%]	13.53	12.04	19.56
Discretionary [Euro]	212.97	139.35	419.58
Non-Discretionary [Euro]	293.53	240.60	469.75
Discretionary / Non-Discretionary [%]	81.46	69.44	102.50
Observations	565,423	37,734	603,157

Notes: This table provides descriptive statistics of characteristics for users who responded to the survey and those who did not respond. The last column reports the standard deviation of the corresponding characteristic across all users. We limit the sample to users that activate the overdraft facility in our main sample.