Goal Setting and Saving in the FinTech Era*

Antonio Gargano†
University of Houston

Alberto G. Rossi‡
Georgetown University

February 12, 2022

We study the effectiveness of soft, self-designed commitment devices, i.e. saving goals, in increasing individuals' savings using data from a FinTech App. We establish that setting goals increases individuals' savings rate and show that the effect is causal using a difference-in-differences identification strategy that exploits the random assignment of users into a group of beta-testers who can set goals and a group of users who cannot. The increased savings within the App do not come at the expense of reduced savings outside the App and that goal-setting also helps the individuals the literature has identified as the ones with the lowest propensity to save. We explore the economic channels of our results by matching App user survey responses to their behavior and highlight the importance of a monitoring channel, consistent with models where agents experience disutility from falling short of their goal and goal setting helping individuals to follow through on their intentions by increasing their attention.

Keywords: FinTech, Goal Setting, Household Finance, Saving Behavior
JEL Classification: D14, G41, G51

*The paper has benefited from comments and suggestions at the 2021 European Finance Association Meeting (EFA), 2021 Boston College Consumer Finance Workshop, the 3rd Future of Financial Information conference, the 2021 annual meeting of the Midwest Finance Association (MFA), the 4th Shanghai-Edinburgh FinTech conference, the 6th FinTech International conference, the 2021 Academic Research Colloquium for Financial Planning and Related Disciplines, and seminar participants at the University of Zurich, the University of Southern California, the University of Missouri and John Hopkins Carey Business School. We are grateful to John Beshears, James Brugler, James Choi, Francesco D'Acunto (the Fintech International conference discussant), Kevin Davis, Marco Giacoletti, Megan Hunter (the BC Consumer Finance Workshop discussant), Claudio Labanca, Olivia Mitchell, Michaela Pagel, Cameron Peng (the EFA discussant), Alessandro Previtero (the Future of Financial Information conference discussant), Juan Sotes-Paladino, Wenlan Qian, Miranda Reiter (the CFP discussant), Donghwa Shin (the MFA discussant), Huang Tang, for comments and suggestions. This research was conducted with restricted access to data from Gimme5, a subsidiary of AcomeA. The views expressed here are those of the authors and do not necessarily reflect the views of Gimme5 or AcomeA. We would like to thank Giuseppe Codazzi and Flavio Talarico from Gimme5 for answering many data related questions and for many insightful conversations. We acknowledge support from the 2020 INQUIRE Europe Research Grant. All errors are our own.

†C.T. Bauer College of Business, University of Houston. Email: agargano@bauer.uh.edu

‡McDonough School of Business, Georgetown University. Fellow of the Luohan Academy. Email: agr60@georgetown.edu.

"Ignoranti quem portum petat nullus suus ventus est"

"If one does not know to which port one is sailing, no wind is favorable"

Lucius Annaeus Seneca

The economics and finance literature has uncovered a number of behavioral biases as the main psychological reasons a large majority of individuals does not save enough to withstand adverse income shocks and is unprepared for retirement.[1] Much less is known about how to correct them on a large scale.[2] In this paper we provide evidence that FinTech Apps can help individuals improve their consumption-saving choices by helping them set saving goals as commitment devices.

Theoretically, commitment devices arise endogenously as cures for undersaving when individuals are affected by behavioral biases such as time-inconsistent preferences (Thaler and Shefrin, 1981; Laibson, 1997; O'Donoghue and Rabin, 1999).[3] Commitment devices vary in terms of whether individuals can deviate from them—*hard* versus *soft*—and whether they are *self-designed* by the users or *pre-designed* by experts. Hard commitment devices generally entail penalties for not following the planned course of action, such as withdrawing funds from a savings account before a pre-established horizon, see Ashraf, Karlan, and Yin (2006). Less is known about soft commitment devices (Beshears, Milkman, and Schwartzstein, 2016), where individuals face no penalties when deviating from their plan. The key trade-off between these two commitment devices is that soft commitments are more scalable and easier to administer compared to hard commitments, especially in non-experimental real life settings. At the same time, the lack of monetary costs associated with non-compliance may make soft commitments less effective (DellaVigna and Malmendier, 2006).[4]

For the most part, the economics literature has shown the effectiveness of pre-designed saving goals (Bryan, Karlan, and Nelson, 2010). The psychological literature, on the other hand, has highlighted

[1]For evidence on undersaving, see, among others, Banks, Blundell, and Tanner (1998); Bernheim, Skinner, and Weinberg (2001); and Lusardi and Mitchell (2007). For the behavioral biases explaining undersaving, see, among others, inertia (Madrian and Shea, 2001), present bias (Laibson, 1997), over-extrapolation (Choi et al., 2009), limited cognitive ability, willpower and self-efficacy (Benartzi and Thaler, 2007; Kuhnen and Melzer, 2018).

[2]See nudges and defaults (Thaler and Sunstein, 2009; Madrian and Shea, 2001; Thaler and Benartzi, 2004; Brown and Previtero, 2020; Medina and Pagel, 2021b), reminders (Karlan et al., 2016b), policy interventions (Gargano and Giacoletti, 2021) and robo-advising (D'Acunto, Prabhala, and Rossi, 2019; Rossi and Utkus, 2020).

[3]For evidence of present bias on debt repayment, see Kuchler and Pagel (2021).

[4]Bénabou and Tirole (2004) and Hsiaw (2013) show that soft commitments can be effective if individuals experience disutility from falling short of their goals. More broadly, Kuhnen and Knutson (2011) show the importance of emotions in financial decisions.

that the effectiveness of goal-setting hinges on the appropriateness and correct design of the goal—see Locke and Latham, 2002 and Locke and Latham, 2006—and it is not clear whether individuals are able to create well-designed saving goals for themselves. Answering this question is crucial given that self-designed goals are more practical and a large number of banks and FinTech firms are now offering goal setting tools for saving to their customers—see Figure A.1 for some examples.

In this paper, we study the effectiveness of *soft, self-designed* commitment devices in saving decisions using data from a FinTech App that provides its users with simple budgeting solutions in the form of saving goals. Our setting is particularly apt at measuring the effect of goal setting uncontaminated by other confounding factors as the App did not include—over our sample period—reminders, notifications, automatic saving plans, or other nudges that have been shown to be effective in promoting saving.[5] To distinguish among the various economic channels underlying the effectiveness of goal-setting, we complement the administrative App data with a survey administered to a sample of App users (Liu et al., Forthcoming).

Because the saving goals we study are self-designed, we start by analyzing their characteristics and assessing the extent to which they are realistic, achievable and congruent with finance theory. First, savings goals are generally short-term, in that the vast majority of the goals have horizons of one year or less. Second, savings goals are generally small, less than €2,500, and we observe evident bunching at whole amounts such as €1,000, €2,000, €5,000, €20,000, etc. Finally, consistent with the view that users understand time value of money and perceive stocks as safer in the longer run, we find a positive relation between goal horizon and both the goal amount and the risk of the chosen investment fund. In both cases, the relation is concave, indicating that neither quantity increases proportionally with the horizon of the goal.

Our central findings pertain to the relation between goal-setting and saving behavior. We start with baseline user-level panel regressions that control for individual-specific and time fixed-effects. Individuals save an additional €345 per year, on average, when they engage in goal setting. Before using the goal-setting feature, the same individuals were saving €345 per year on the App, suggesting that goal setting doubles individuals' savings. Economically, these effects are remarkable, given that—

[5] Our sample ends in November 2019. Some of these features were introduced over the course of 2020, others are being launched in the Fall of 2021.

as we stress above—the App does not include reminders, notifications, automatic saving plans, or other saving nudges. When we decompose the overall savings into its parts, i.e., withdrawals and deposits, we find the additional savings are mainly driven by an increase in the deposit rate rather than a decrease in the withdrawal rate.

The limitation of the user-level results is that individual-specific time-varying shocks could explain the differential savings when individuals set goals. In this respect, our setting is unique in that we can compute results at the user-account-level, which allows us to compare the saving behavior of a user who sets goals to the saving behavior of the same user in other accounts where she does not set goals, controlling for individual time-varying shocks using user×time fixed effects. This empirical strategy suggests an even stronger effect of goal setting on saving of €661.9 per year. In line with the user-level results, the user-account-level results also show that goal setting increases individuals' propensity to deposit and leaves users' withdrawal activity unchanged.

We augment these results with dynamic specifications that study the effectiveness of goal setting over time. Users save substantially more in the first month after adopting goal-setting, where additional savings exceed €65 per month. Over time, however, the additional savings stabilize to approximately €30-35 per month. This indicates that the initial enthusiasm induced by the new saving tool explains part, but not the full extent, of the savings differential we document in the baseline results. Moreover, we show that goal-setting increases the consistency with which individuals deposit money (extensive margin) on the saving App rather than the sums of money they deposit every time they decide to save (intensive margin).

Our setting also allows us to identify causal effects using a quasi-natural experiment. Ideally, one would want to implement a Randomized Control Trial (RCT) where some individuals are randomly assigned to be goal-setters and others are randomly assigned to be non-goal-setters for an extended period, i.e. several months or years. Fortunately, while we do not have access to such randomized experiment, we do have access to something very similar. The App developers activated the goal-setting feature for 122 individuals (beta testers) 50 days before the goal-setting feature was rolled-out to the whole population. This allows us to estimate a difference-in-differences model where the beta testers are the treated group and the rest of the users are the control group. Because the users in the treated group had access to the goal-setting feature, but were not forced to use it, with this identification strategy we can estimate an Intention to Treat (ITT) effect—the effect of deploying goal setting on the broader population at large—as well as a Local Average Treatment Effect (LATE)—the effect of goal setting for those who actually set goals. We estimate an ITT of €203.4 per year and a LATE of €714.8 per year, confirming our baseline results.

To rule out that the increased savings we observe inside the App are a substitute for the savings outside the App, we explicitly ask the App users in our survey how they have changed their saving habits outside the App after its adoption and show that, if anything, they have increased their saving outside the App. We also re-estimate our baseline findings on those individuals who declare they do not save outside the App and show that the results are very similar to the ones we find on the overall population.

In the third part of the paper we explore the economic channels that can explain the effectiveness of goal setting (Beshears, Milkman, and Schwartzstein, 2016). We first consider a monitoring channel whereby goal-setting increases savings because individuals experience disutility from falling short of their goal (Hsiaw, 2013; Bénabou and Tirole, 2004) and/or goals help individuals to follow through on their intentions by increasing their attention (Karlan et al., 2016a). We also explore a goal concreteness channel whereby the specific purpose of the goal creates a specific segregated mental account that acts as a saving motivator (Thaler, 1985, 1990, 1999; Shefrin and Thaler, 1988). A large psychology literature also establishes that setting explicit and concrete goals increases achievement, even in the absence of external incentives Locke and Latham (1990).

While we cannot fully disentangle between the different subchannels discussed above, we show that monitoring is the main driver of goal effectiveness using three different empirical exercises. First, we show that, when surveyed, individuals identify the ability to track their progress towards a goal as the most motivating aspect of setting a goal. Other aspects such as the purpose of the goal or the presence of a deadline are less important.

Second, we relate users' saving behavior to the characteristics of the goals they set such as type, horizon, and goal amount. When we group goals based on type (*leisure*, which includes Hobby and Travel; *durable*, which includes Car and House; and *generic* saving) we find no significant differences, suggesting that the purpose of the goal does not matter. We also find individuals save significantly more for shorter goals (less than or equal to one year), consistent with the shorter horizon making it easier for individuals to monitor their progress.

Third, we shed light on the factors affecting the probability of achieving goals as well as the dynamics of saving for goals. While goal-setting increases savings, it does not imply that users reach their goals. Indeed, we find that only 30% of the goals are achieved, while the rest are either closed or remain overdue after their deadline has passed. The specificity of the goal does not increase the completion rate. On the contrary, specific goals such as hobby, travel, car and house categories all have a probability of success 10% lower than the "generic savings" category.[6] We also find that short-term and small goals are the ones most likely to be achieved. Taken together these results reinforce the idea that the increase in savings we document in our main results is likely not due to the purpose of the goal, but rather to users' ability to monitor their progress towards their goal.

In the final section of the paper, we study whether goal-setting can help those individuals the literature has found to be particularly at risk of under-saving and whether the increases in savings associated with goal-setting are economically large relative to users' income. The most common reasons behind individuals' inability to save range from low levels of financial literacy (Lusardi, 2008 and Lusardi and Mitchell, 2014), high impatience (Laibson, 1997), low levels of overall education (Cole, Paulson, and Shastry, 2014 and Bernheim, Garrett, and Maki, 2001), low income and socioeconomic status (Banerjee and Mullainathan, 2010, Duflo and Banerjee, 2011 and Kuhnen and Miu, 2017), and user's inattention (Karlan et al., 2016b). Our results suggest that even those who are more prone to under-saving increase significantly their saving after engaging in goal-setting. If anything, in the majority of instances we find that the least and most vulnerable users benefit equally from goal setting.

To understand the economic magnitude of goal-setting on users' savings rate, we repeat the analysis on the subset of survey respondents for which we have income information. We use three quantities: 1) individuals' saving rate before adopting goal setting; 2) the implied saving rate associated with saving goals; and 3) the saving rate users achieve when they use goal-setting, which is the result of planning and having the freedom to deviate from the original plan. The difference between the saving rate

[6] To make sure our findings are not an artifact of the coarse categories set by the App, in a robustness section we confirm these results using categories based on user-generated goal names.

implied by goals and the saving rate when not setting goals can be thought to represent the potential effect of goal setting in increasing savings rates. The difference between the saving rate implied by the goals and the saving rate obtained with goal-setting can instead be taken to represent the reduction in saving (leakage) individuals experience with soft commitments. Finally, the difference between the saving rate with goals and the saving rate without goals measures the realized effect of goal setting.

The median individual saves 1% of her net income when she does not set goals. The planned percentage savings associated with goals equal 4.6% of her net income. Finally, the percentage savings realized when setting goals equal 2.4% of her income. These differences in savings rates highlight that there is substantial leakage associated with the soft commitments individuals engage in. They also show that goal setting can help individuals significantly increase their baseline savings rate.

1 Related Literature

This paper contributes to the very recent and fast-growing literature on Financial Technology (FinTech) and household welfare. FinTech encompasses a broad range of new technologies that seek to improve and automate the delivery and use of financial services (see Das, 2019 for a review of the literature). A recent strand of the literature shows that FinTech can help households improve their trading performance and asset allocation (e.g. Gargano and Rossi, 2018; D'Acunto, Prabhala, and Rossi, 2019; Rossi and Utkus, 2020; and D'Acunto and Rossi, Forthcoming), to reduce overspending (D'Acunto, Rossi, and Weber, 2019), to save on bank fees (Carlin, Olafsson, and Pagel, 2019; D'Acunto et al., 2019; and Loh and Choi, 2020) and to better conduct house searches (Gargano, Giacoletti, and Jarnecic, 2019).[7] At the same time, another strand of the literature highlights the pitfalls generated by the introduction of new technology. For example, Fuster et al. (2018) find that Black and Hispanic borrowers are disproportionately less likely to gain from the introduction of machine learning tools to predict creditworthiness, while Di Maggio and Yao (2020) find that FinTech borrowers are significantly more likely to default than their peers borrowing from traditional financial institutions. Laudenbach, Pirschel, and Siegel (2018) find that borrowers who speak directly with a bank agent are significantly less likely to default. We contribute to this literature by showing the benefits of financial technology

[7] More broadly, Agarwal et al. (2020) and Agarwal, Yeung, and Zou (2019) show the positive effect of Financial Technology on entrepreneurship and business creation.

in the savings domain, which still remains largely unexplored.

We also contribute to the literature that studies commitment devices. The effectiveness of hard commitments, which either entail a monetary penalty for deviating from the plan of action or restrict the set of future available options, is supported by evidence across a broad range of domains. For example, these commitment devices have been shown to increase agricultural input use (Duflo, Kremer, and Robinson, 2011), preventative health investment (Dupas and Robinson, 2013), and chances of successful smoking cessation (Giné, Karlan, and Zinman, 2010). In the saving domain, the focus of our study, Ashraf, Karlan, and Yin (2006) run an experiment in the Philippines and show that individuals with higher hyperbolic discounting are more likely to open a savings product restricting access to deposits which in turn increases their saving. Beshaers et al. (2020) study the effectiveness of different schemes and find that higher early-withdrawal penalties attract more commitment account deposits. On the other hand, the evidence on soft saving commitments, where individuals face no penalties when deviating from their plan, is more limited. These commitments are more scalable and much simpler to administer compared to hard commitments, but the lack of monetary costs associated with non-compliance may make soft commitments less effective (DellaVigna and Malmendier, 2006). Moreover, it is not clear whether individuals are able to create well-designed saving goals for themselves. Breza and Chandrasehkar (2019) run an experiment in rural India and show that individuals save more when information about the progress toward their self-set savings goal is shared with another village member who act as a monitor. On the other hand, John (2020) shows that imperfect knowledge about their preferences leads individuals to select into incentive-incompatible saving commitments, which in turn reduce their welfare. We contribute to this literature by providing evidence on *soft*, *self-designed* commitments that can administered on a large scale.

2 The Setting

This section provides a description of the App and the survey we conducted on a random sample of its users.

2.1 The Saving App and the Introduction of the Goal-Setting Features

The data used in this study were shared by the Italian asset manager AcomeA that developed a saving App called Gimme5. Launched in 2014, it was initially designed as a digital piggy-bank that allowed its users to save small sums of money, such as 5 euros, at a time—hence the name "Gimme5." At the time of sign-up, individuals face zero activation costs and no annual fees.[8] After signing up, users can start saving by depositing as much money as they wish, whenever they want, much like a checking/savings account. Unlike checking/savings accounts, however, Gimme5 users invest their savings—however small they are—in several investment funds that vary in their risk-return profiles. As of 2019, individuals could choose from as many as 14 funds, some of which are actively managed and therefore charge management fees. Table A.1 displays the annualized return, volatility and Sharpe ratio of each fund computed over the Jan-2015 to Dec-2019 period.

In September 2017, the App introduced a goal-setting feature through which users can set one or multiple saving goals. Since then, when setting a goal, users are prompted to choose the objective, categorized in six broad categories: hobby, travel, other, car, house and general savings. The users are also required to choose the horizon and the target amount. Figure 1 presents screenshots of the App over these phases. The first step (Subfigure A) prompts the user to name the target. The second (Subfigure B) asks the user to choose between the six categories: Hobby (Tempo Libero), Travel (Viaggi), Other (Altro), Car (Veicoli), House (Casa) and Savings (Risparmio). The third (Subfigure C) prompts the user to choose a target amount and a deadline to reach the target. Finally, the fourth (Subfigure D) asks the user to associate the target with an image. The App is by no means unique or specific to the Italian setting. In recent years, many banks like Chase, Wells Fargo and financial aggregators in the US such as MINT.com have implemented saving goal tools that are virtually identical in structure and features as the one we describe here—see Figure A.1 for some examples.

After setting up the target, the user chooses which mutual fund she would like to invest in. The App features three baseline funds, displayed in Figure 2 (Subfigure A). Their names are Prudent (Prudente), Dynamic (Dinamico), and Aggressive (Aggressivo), and they are characterized by increasing levels of risk. The users can choose among 14 total funds, by clicking on the link "Full List of Funds" (Elenco

[8] The only fee they face relates to the withdrawal of funds. Users are charged a €1 withdrawal fee, irrespective of the amount withdrawn.

Completo dei Fondi).

After the goal and the fund are chosen, the user can deposit money whenever she likes. The App features an interface that shows the user where she stands in achieving the goal. An example is reported in the two screenshots—Subfigures B and C—of Figure 2. In this example, the user has budgeted €3,000 for a vacation in the US and has 568 days to meet the goal. As of the day of the screenshot, the person has saved €10. The graphics makes salient how far the user is from the target. The user can click on the pig icon at the bottom of the screen—referred to as "Joink"—to add funds to the account and reach the goal. Once the user pushes the Joink button, she is taken to a different page where she can choose how much to deposit for that goal, the default being €5.

Crucially, after the introduction of these goal-setting features, users could choose to continue using the App as a digital piggy-bank with no goals, use only the newly introduced goals, or use both features at the same time.

As stressed in the introduction, an important feature that allows us to identify the effect of goal-setting on saving is that at no time during our sample the App had reminders, notifications, automatic saving plans, or other saving nudges. Some of these features were introduced in 2020, while others are being launched in the Fall of 2021.

2.2 The Survey

We complement the administrative data with a survey administered to a sample of App users.[9] Matching individuals' actions on the App to their survey responses is crucial to uncover the economic mechanism underlying individuals' choices, see Liu et al. (Forthcoming).[10] The survey comprises four sections. In the first, we elicit information we cannot observe from the administrative data. For example, we ask App users how much they were saving outside the App before and after its adoption. We also ask them what are the main accounts they use to save and invest their wealth, that is, savings accounts, brokerage accounts, retirement account, etc.

[9]We report the translated survey in the the last section of the Online Appendix.

[10]For additional examples, see D'Acunto, Malmendier, and Weber (Forthcoming), Coibion et al. (2020), Coibion, Gorodnichenko, and Weber (2019), Choi and Robertson (2020), Brauer, Hackethal, and Hanspal (2020), Giglio et al. (Forthcoming), and Giglio et al. (2021). The survey was administered by email to a random sample of 5,000 App users over the week of March 15-21, 2021. The number of respondents was 814, for a response rate of 16.3%.

The second and third sections of the survey focus on individuals' motivations for creating targets within the App and using the Joink feature. More precisely, for those who have created at least one target, we ask what they believe is the most motivating aspect of goal-setting, and what are the main reasons they did (or did not) achieve a target in the past. Finally, we elicit—in a number of questions—what is the underlying psychological motivation for using the Joink feature.

In the last part of the survey we measure users' patience, risk-aversion, and financial literacy following standard questions in the literature—see Lusardi and Mitchell (2017), Falk et al. (2018), and D'Acunto, Rossi, and Weber (2019). We also ask respondents for their level of education, and net income. Figure A.2 reports a summary of the survey responses regarding financial literacy, patience, risk propensity, education and monthly income.

3 Data and Self-designed Goal Characteristics

Next we describe the main data sources and report the summary statistics for the key variables of interest. All information was anonymized by the company to guarantee user privacy.

3.1 Data Sources

The data is in the form of 5 SQL tables named *Users*, *Goal Information*, *Deposits and Withdrawals*, *Investment Vehicle*, and *Login Information*.

Users. This table contains information on the 47,643 users who created a profile on the App since its inception. The main variables contained in this table are the dates of opening and closing for every savings account created by each user.[11] Additional information includes users' age, gender, and enrollment date—the date on which the profile was created—and the locations where the users reside, where they where born, and other questions related to Know Your Customer (KYC) compliance.

Goal Information. This table contains information on all the goals set by each user in each savings account. For each goal we observe its inception, its deadline, its amount, its purpose—categorized in the six broad categories: hobby, holiday, other, car, house, and "generic savings"—and its name.

Deposits and Withdrawals. This table contains information on all the deposits and withdrawals

[11] Users can create multiple savings accounts within the App to invest their savings across different mutual funds.

associated with each user at the daily frequency. For each transaction, we have information on which account and/or goal it is associated with, as well as the amount withdrawn or deposited.

Investment Vehicle. This table contains information on the name and identifier of the mutual fund associated with each savings account of the user.

Login Information. This table contains login time-stamp information for every App user.

3.2 Summary Statistics

Table 1 reports the main cross-sectional summary statistics, computed in two steps. We first compute the value of each variable at the user level and then report the distribution of the variable across all users. For each variable, we report the number of observations used in the second step of the computations, the mean and standard deviation of each variable, as well as the $1st$, $25th$, $50th$, $75th$, and $99th$ percentiles.

Panel A shows that 81% of the users are males, suggesting that, in Italy, women are either less targeted by Financial Apps or are less interested in these tools. Users are rather young, with an average age of 36. Panel B reports results on App usage. Here the summary statistics are computed using individuals who have transacted at least once on the platform, which are 27,439. As of November 2019—the date when we received the data— the average user had used the App for 11.50 months. Active users interact with the App often and deposit €30.04 per month.

Panel C reports facts regarding the number of accounts and targets associated with each user. Turning to targets, 27 percent of the users have created savings goals as of November 2019, indicating that the majority of users used the App as a simple "Digital Piggy Bank" without setting specific goals for their savings. Individuals like to set one or—at most—two targets at a time (the average is 1.43). The average investor had a little less than one target open and the average number of targets closed and not completed was equal to 0.3. Finally, the number of targets completed equaled to 0.28 per user.

3.3 Saving Goals

Because the saving goals we study are self-designed, their analysis can offer a window into users' savings preferences and beliefs, and allows us to assess the extent to which they are realistic, achievable and congruent with finance theory.

The first two panels of Figure 3 report the histograms of target horizons. To improve the legibility of the graph, we split the histogram into two plots, the first associated with horizons shorter than 1,000 days, and the second associated with horizons longer than 1,000 days. Overall, individuals set savings goals that are rather short-lived. The majority of the goals have horizons of less than one year and a large number of goals have horizons of exactly one year. Furthermore, we observe clear bunching of the distribution at the one, two, three, all the way to the 10-year horizon.

The second two panels of Figure 3 report histograms of goal amounts. Once again, individuals set goal amounts at round numbers. The majority of the goals are small, less than €2,500. For larger amounts, the distribution has evident bunching at €1,000, €2,000, €3,000, €5,000, €10,000, €20,000, and €50,000.

The bottom-left panel of Figure 3 focuses on the 14 mutual funds from which investors can choose to invest their savings. We classify the funds in three categories based on their type: Fixed-income, Balanced and Equity—see Table A.1 for details. As shown in the bottom-left panel, the majority of investors (55%) chooses balanced funds, followed by fixed-income and equity funds that are selected by 32% and 13% of the users, respectively. Individuals avoid exposure to pure equity funds, suggesting a relatively high degree of risk-aversion and low trust in the stock market (Guiso, Sapienza, and Zingales, 2008).

The bottom-right panel of Figure 3 focuses on goal categories. Half of the individuals set goals that do not fall naturally into specific categories and are therefore assigned a generic category named "Savings," suggesting that the purpose of the goal may not be crucial to the App users—an aspect we study at great length in Section 5. Turning to the other categories, "Travel" is the most common, which comprises 20% of the targets, followed by "Car," "Hobby" and "House," which each make up a little less than 10% of the categories.

Next, we relate the horizon of each target to the target amount and the investment risk using a nearest neighbor estimator in Figure 4. We find a positive and monotonic relation between target horizon and amount (Panel A). Quantitatively, users budget approximately 1,000 days to save €5,000; 2,000 days for €10,000. The relation then becomes more concave, indicating that individuals budget proportionally more time to save additional amounts. We also find an increasing and concave relation between target horizon and investment risk (Panel B). As the investment horizon increases users choose riskier funds, supporting the idea they believe stocks are a better investment for the long-run compared to fixed-income securities (Siegel, 1998).

4 The Effects of Goal-Setting on Saving Behavior

Results from the psychology literature (Locke and Latham, 2002, 2006) suggest that goal-setting should have a positive effect on saving since setting goals allows individuals to remain focused on the saving objective and deadline, and allows them to track their performance against it. In addition, several theoretical models developed in economics (Hsiaw, 2013; Bénabou and Tirole, 2004; Besshears, Milkman, and Schwartzstein, 2016) predict that present-biased individuals set goals as a soft commitment device to overcome their self-control problems. Such goals provide internal motivation because individuals derive utility from comparing outcomes to goals.

On the other hand, goals could be counterproductive and have a discouraging effect if they are too ambitious and unrealistic—see, for example, John (2020). This concern is particularly relevant in our setting as individuals choose their goal characteristics, such as horizon and amount and the goals we study are non-binding commitments.

4.1 Results at the User Level

In our baseline specification, we estimate a panel regression of the following form:

$$M\_Net\_Deposits_{i,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,t} + \epsilon_{i,t}, \quad (1)$$

where $M\_Net\_Deposits_{i,t}$ are the monthly net deposits (deposits minus withdrawals) of individual i during month t, $Target\_Dummy_{i,t}$ is an indicator variable equal to "0" if the user did not set any goals in her account at time t and "1" if she set at least one goal at time t. The coefficients α_i and α_t denote individual and time effects, respectively. The standard errors are double-clustered at the user and month levels.

The coefficient of interest is β, which identifies the change in monthly savings when users engage in goal-setting, controlling for individual-specific characteristics and time-effects. The results, reported in the first column of Table 2, indicate a large effect of goal-setting on saving behavior. The coefficient equals 28.74, with a t-statistic of 16.37, indicating that individuals save an additional 28.74×12=€345 per year when they set goals. The average saving across all users when not setting goals equals 28.75, or €345 per year. Thus users double their savings when they set goals. Economically, these effects are remarkable, given that the App does not include reminders, notifications, automatic saving plans, or other saving nudges.

In the second and third columns of Table 2 we repeat the analysis on users' deposits and withdrawals, respectively. The coefficient on the $Target\_Dummy_{i,t}$ suggests that those who set goals save an additional €34.56 when setting goals, compared to when they don't. This quantity is economically large, given that average deposits equal €36.43 per month when individuals do not set goals. The effect of goal-setting is less pronounced for withdrawals. Goal-setters withdraw an additional €5.60 compared to when they do not set goals, but the economic magnitude of the effect is rather small. By comparing the results for Net Deposits, Deposits, and Withdrawals, it is clear that goal-setters deposited more in their accounts and also withdrew more. The additional withdrawals, however, are only a fraction of the additional deposits. As a result, the change in net deposits is strongly positive and significant.

4.2 Results at the User-Account Level

One limitation of the results reported above is that time-varying individual-specific shocks could explain the differential saving behavior of individuals when they set goals compared to when they don't. Our setup is unique in this respect, in that it allows us to estimate results at the user-account level, that is, we can compare the behavior of a user who sets goals in one (or more) of her savings accounts, to the behavior of the *same* user in other savings accounts where she did not set goals. We compute these estimates using the panel regression:

$$M\_Net\_Deposits_{i,j,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,j,t} + \epsilon_{i,j,t}, \tag{2}$$

where $M\_Net\_Deposits_{i,j,t}$ are the net deposits of user i in account j at time t; $Target\_Dummy_{i,j,t}$ is dummy variable equal to 1 if account j associated with user i at time t is characterized by goal-setting, and zero otherwise. The coefficient of interest (β) measures the differential saving between accounts with targets and accounts without targets.

The results of Specification 2 are reported in Panel A of Table 3. In the first column, the coefficient estimate on $Target\_Dummy_{i,j,t}$ equals €26.84, which implies a differential saving of €26.84 $\times$ 12=€322.08 per year in accounts with targets compared to accounts without targets, controlling for time-invariant user characteristics. The coefficient for Deposits (column 2) is also large (€29.40) and has a t-statistics in excess of 15. The one for Withdrawals (column 3) is instead small and insignificantly different from zero.

In Panel B, we repeat the analysis, but include user-time fixed effects to control for time-varying individual characteristics that may be driving changes in saving, that is, we estimate:

$$M\_Net\_Deposits_{i,j,t} = \alpha_{i,t} + \beta\ Target\_Dummy_{i,j,t} + \epsilon_{i,j,t}, \tag{3}$$

where all the variables are defined as in Specification 2, but the coefficient of interest (β) now measures the differential saving between accounts with targets and accounts without targets when investors use both types of accounts at the same time. Even for this more stringent specification, the coefficient estimates are strongly positive and significant. In the first column, the coefficient estimate on $Target\_Dummy_{i,j,t}$ equals 55.16, implying a differential saving of 55.16 $\times$ 12=€661.92 per year in accounts with targets compared to accounts without targets. As in Panel A, in the second and third columns we find that setting goals increases individuals' propensity to save, but does not increase users' withdrawal activity.

4.3 Goal-Setting and Persistence of Saving Behavior Over Time

The results reported so far could potentially be explained by a "new toy" effect, whereby individuals saved more at the beginning, but stopped saving as they grew accustomed to the novelty of goal setting. We can test how individuals save over time as they adopt the new goal-setting features introduced in the App by estimating a modified version of Equation 3, as reported below:

$$M\_Net\_Deposits_{i,j,t} = \alpha_{i,t} + \sum_{s=-6}^{18} \beta_s \; Target\_Dummy_{i,j,s,t} + \epsilon_{i,j,t}. \tag{4}$$

The only difference between Equation (4) and Equation (3) is in the definition of $Target\_ Dummy_{i,j,s,t}$. In Specification 4, $Target\_Dummy_{i,j,s,t}$ equals 1 only if individual i used goal-setting in account j for the $s-th$ month in the calendar month t. In the specification above, s ranges from -6 to 18. The β_s coefficients in Specification 4 therefore decompose the effect of β in Specification 3.

The coefficient estimates for Net Deposits, Deposits and Withdrawals are reported in Figure 5. The results suggest two findings. First, the effect is strongest over the first six months, where it peaks at €65 per month; and it then decreases and asymptotes to approximately €30-€35 per month, suggesting that initial enthusiasm explains part of the increase in savings, but not the full saving differential. Second, the effect seems to be fully driven by deposits: none of the effect is driven by withdrawals that never differ from zero at any horizon.

The results reported in Figure 5 mix the extensive and the intensive margins. At one extreme, it could be that individuals deposit money in the App more frequently when setting goals, compared to not setting goals. At the other extreme, it could be that individuals deposit larger sums of money every time they make a deposit in their accounts with goals compared to the accounts without goals. Both hypotheses, as well as the intermediate cases, are consistent with the results in Figure 5.

To address whether the effect is mainly driven by the intensive margin or the extensive margin, we estimate two variations of Equation (4). In the first, we replace $M\_Net\_Deposits_{i,j,t}$ with $IND_{i,j,t}$, an indicator variable equal to 1 if the monthly net deposit of individual i in account j is positive in month t. In the second we replace $M\_Net\_Deposits_{i,j,t}$ with $M\_Net\_Deposits^{+}_{i,j,t}$, which denotes the monthly net deposit (in €) by user i in account j at time t, conditional on $IND_{i,j,t}$ being equal to 1.

The β coefficients in the first specification measure the increase in probability of having positive net deposits, while the β coefficients in the second specification measure the increase in amount deposited, conditional on the user depositing wealth in the App on a given month. The results, reported in Figure 6, highlight that goal-setting increases mainly the extensive margin of saving. The probability of depositing money on the App increases by 65% on the first month after setting the goal, but the effect is reduced as the horizon increases. The effect stabilizes to approximately 50%, 12 months after the goal is set. The intensive margin results, on the other hand, show very little changes after goal-setting takes place. Conditional on depositing money on the App, users increase their net deposits by €10 the first and second months after setting the goal, but the increase becomes insignificant from the third month onwards. Overall, these results indicate that goal-setting increases the consistency with which individuals deposit money on the saving App rather than the sums of money they deposit every time they decide to save.

4.4 Identification Strategy Results

Our setting also allows us to identify causal effects using a quasi-natural experiment. Ideally, one would want to implement a Randomized Control Trial (RCT) where some individuals are randomly assigned to be goal-setters and others are randomly assigned to be non-goal-setters for an extended period, i.e. several months or years. Fortunately, the App developers implemented something very similar to a randomized control trial: they activated the goal-setting feature for 122 individuals 50 days before they made it available to the overall population to ensure the reliability of the new version of the App. These selected users were unaware of being selected for the trial and were not volunteers. Consistent with this, only 28.4% of the beta testers end up setting goals in the beta-testing period. The main difference between the procedure implemented by the firm and a standard randomized control trial is that the number of individuals selected for the trial is probably smaller than we would have preferred, and that the company did not store the details of the formal randomization scheme they used. For example, we do not know if they stratified the sample. The 50-day trial duration is another limitation.

Nevertheless, this natural experiment allows us to estimate a difference-in-differences model where the beta testers are the treated group and the remaining users are the control group. Our specification reads as follows:

$$M\_Net\_Deposits_{i,t} = \alpha_i + \alpha_t + \beta \ Treated_{i,t} + \epsilon_{i,t}, \tag{5}$$

where all quantities are defined as in Equation (2) except for $Treated_{i,t}$ which is equal to 1 for the 122 beta testers for the 50 days (from August 1, 2017 to September 20, 2017) over which they had access to goal-setting while the rest did not.

Note that, because we estimate our specification at the monthly level and we seek to be as comprehensive as possible, we estimate two specifications. In the first, the $Treated_{i,t}$ dummy is equal to 1 only for the month of August 2017. In the second, the $Treated_{i,t}$ dummy is equal to 1 for both August 2017 and September 2017. In both cases, we estimate our results using an event window that starts in August 2016. Having a full year of deposit observations before the treatment allays seasonality concerns in our estimated coefficients.

Note also that, unlike the results estimated in Section 4.2, here the β coefficient identifies an Intent-To-Treat (ITT) effect, because the beta-testers are given access to the goal-setting technology, but they are not required to use it. The coefficient therefore estimates whether goal-setting would increase investors' saving over time, if it was deployed on the broader population at large.

The results for net deposits are reported in Panel A of Table 4. Irrespective of how the treatment variable is defined—August 2017 or August and September 2017—we find positive and statistically significant estimates, that is, those who had access to the goal-setting features of the App saved, on average, an additional €16.92-€20.01 per month, or €203.4-240.12 per year, relative to their non-beta-testers counterparts.

The estimates reported above represent ITT effects. To obtain the Local Average Treatment Effect (LATE), we divide the ITT by the percentage of beta-testers that engaged in goal-setting during the beta-testing period. For Specification (2), the percentage of compliers is 28.4%, giving a LATE effect of €16.92/28.4%=€59.57 per month, which is very close to the baseline estimates we obtain in Panel B of Table 3.

As mentioned above, the company did not store the details of the randomization algorithm they implemented to select the beta testers. It is therefore not guaranteed that the beta testers' activity and demographic characteristics match the ones of the rest of the users. We analyze this in Table 5. In order to capture users' activity, we use the following variables: $Tenure$, the number of months elapsed from the opening of user's first account to July 2017—right before the beginning of the beta-testing period; *N. Accounts* and $Net\ Inflow$, which denote the number of accounts opened and the monthly average net deposit into the App, computed from August 2016 to July 2017. To capture user risk-preferences and demographic characteristics, we use: *Risk*, the average annual standard deviation across the funds used by the user; *Gender*, a dummy variable equal to one if male; and *Age*, measured as of July 2017. Panel A displays the characteristics of the beta testers while Panel B focuses on the rest of the users in the sample. For this second sample, we also report the t-statistics associated with the tests of the equality of means between beta-testers and non-beta-testers. The results indicate that the two groups of users are rather similar. At the 5% significance level, beta-testers and non-beta-testers are indistinguishable in terms of all the characteristics, except for $Tenure$. Beta testers are statistically more experienced in that they have been on the App for longer (516 months against 431). In order to account for such differences and create a sample of users more comparable with the beta testers, we use a nearest neighbor propensity score matching procedure. For each beta-tester, we choose the closest match based on the six variables described above.[12] Panel C shows the characteristics of the matched sample of users. All the t-tests are insignificant, indicating that these matched users are indeed statistically indistinguishable from the beta testers.

Even though the results in the Table 5 do not highlight major differences between beta-testers and non-beta-testers, in Panel B of Table 4 we re-estimate the difference-in-differences specification in Equation (5) using the matched sample described above as well as when matching each beta tester with the closest two or five users. The results highlight that, even though the sample has dramatically dropped in size, the coefficients remain statistically significant and are very similar to the ones obtained without matching, allaying concerns regarding the randomization procedure implemented by the company to select the beta testers.

The strength of the identification results reported here are further confirmation of the effectiveness of goal setting on saving.

[12] In Table A.2 we report results when matching each beta tester with the closest two and five users.

4.5 Ruling Out Substitution Between Saving Within and Outside the App.

A possible concern with our results is that individuals could perceive saving within and outside the App as substitutes, and could therefore decrease their savings outside the App when they start using the App. Unfortunately, we cannot rule out this mechanism with administrative data only, as we do not have information on deposits and withdrawals on accounts other than the ones on the App.

To circumvent this limitation, we included in our survey specific questions related to individuals' saving outside the App, before and after its adoption. In particular, we ask three questions. The first asks users whether they increased, decreased or left unchanged their savings rate outside the App after its adoption. The results are reported in Panel a) of Figure 7. The vast majority of the respondents (51.2%) say they have not changed their savings outside the App. A significant number (35.0%) state they increased their saving outside of the App. Finally, only a small minority of individuals (13.8%) state they decreased their savings outside the App.

In two additional survey questions, we explicitly ask the respondents for their saving rate before and after adopting the App and gave them five options to choose from: 1) €0; 2) €1-€49; 3) €50-€100; 4) €101-€150; 5) €150+. As shown in Panel b) of Figure 7, individuals increased their savings outside the App after its adoption. The red bars—denoting saving after adoption—are higher than the blue bars—denoting saving before adoption—for all response categories except for the €0 option. This entails that the entire distribution of saving outside the App increases after individuals sign up, ruling out any potential substitution effect between saving inside and outside the App.

We also perform several additional quantitative tests to rule the possibility of a substitution effect between saving inside and outside the App after adopting goal-setting. We report them in Table A.3. In the first column, we re-estimate regression specification (1) only on the survey respondents and find a very similar effect as on the full sample (column 1 of Table 2). In the second column, we restrict the sample to those survey respondents who state they have no savings outside the App and show the results are virtually identical.

In the last three specifications, we interact the target adoption dummy with three groups of users. The first comprises users who do not save outside the App. For these users the substitution channel cannot be present by construction, we therefore expect a lower increase in saving after adopting the target (negative interaction coefficient) if a substitution effect was in place among the rest of the App users. The second group comprises users who state they do not change their saving rate outside the App after its adoption. Also for these users, the substitution effect should be zero and we should expect a negative interaction coefficient. Finally, the third comprises users who state they reduce their savings outside the App after its adoption. These users should experience the maximum substitution effect, compared to the rest of the users. We should therefore expect a positive and significant interaction coefficient. In all cases, we find that the coefficients on the interactions are economically small and not statistically significant, suggesting there is no evidence of substitution between saving within and outside the App on our data.

An additional concern could be that individuals increased their debt levels—particularly credit card debt—after increasing their savings on the App as a result of goal setting. This is unlikely to occur for two reasons. First, recent studies show that individuals who increase savings as a result of saving nudges do not increase their debt balances subsequently (Besheears et al., Forthcoming; Medina and Pagel, 2021a). Second, the use of credit cards is much less widespread in Italy, compared to the US. According to the Euromonitor International Consumer Finance 2019 report, only 26% of the households hold a credit card in Italy. The corresponding percentages for the rest of Western Europe and the US are, respectively, 45% and 69%.

5 Why Does Goal-Setting Increase Savings?

The psychology and economics literatures show the two most important aspects of setting goals rests on the concreteness of the goal and users' ability to monitor their progress (Beshears, Milkman, and Schwartzstein, 2016). Goal concreteness can inspire higher savings due to mental accounting, because the specific purpose of the goal creates a specific segregated mental account that acts as a saving motivator (Thaler, 1985, 1990, 1999; Shefrin and Thaler, 1988). A large literature in psychology also demonstrates that setting explicit and concrete goals increases achievement, even in the absence of external incentives (Locke and Latham, 1990). The monitoring channel instead implies that savings should increase with goal setting because individuals experience disutility from falling short of their goal (Hsiaw, 2013; Bénabou and Tirole, 2004) and/or goals help individuals to follow through on their intentions by increasing their attention (Karlan et al., 2016a).

While we cannot fully disentangle between the different subchannels discussed above, in this section we show that monitoring is the main driver underlying the effectiveness of goal setting for saving using three different empirical exercises. First we show that, when surveyed, individuals identify monitoring as the most important channel rather than the specific purpose of the goal (Section 5.1). Second, we study how the type of goal (Section 5.2) relates to saving behavior, and show that more specific goals are not associated with higher saving rates. Third, we study the main determinants of goal achievement (Section 5.3), and show that also for this dimension the specific purpose of the goal does not matter.

5.1 Survey Evidence

We start with reporting in Panel a) of Figure 8 the distribution of the answers to the question *"In your experience, what is the most motivating aspect of having a saving goal?"* The options we provided emphasized three distinct aspects of goals. The first two entailed the deadline dimension and the ability to track the progress towards goals. The third instead related to concreteness and entailed the purpose of the goal. We provided users with the following choices: 1) Having a deadline; 2) The ability to monitor my progress towards the goal; 3) The purpose of the goal; 4) Other; and 5) I do not think having a goal is motivating.

App users overwhelmingly respond (57%) that the ability to monitor their progress towards the goal was the most motivating aspect, followed by the purpose of the goal (28.5%) and having a deadline (10.4%).

We corroborate these results using observational data on individuals' login activity on the App, before and after setting goals. We estimate a regression identical to specification (1), except that we replace the monthly net deposit dependent variable with the (logged) number of monthly logins on the App. We find that individuals increase their login activity on the App by 6% on average, consistent with the intuition that goal setting increases the amount of time and energy people spend thinking about their savings.

5.2 Specificity of Goal and Saving Behavior

In Table 6, we re-estimate the models in Table 2, but condition our estimates on the type of target. The first column splits the type of targets into four categories: *leisure*, which comprises Hobby and Travel; *durable*, which comprises Car and House; *other*, collecting other specific targets; and *savings*, the generic saving category. We then estimate a regression model:

$$M\_Net\_Deposits_{i,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,t} + \sum_{k=1}^{3} \gamma_k\ Target\_Dummy_{i,t} \times Interaction\_k_i + \epsilon_{i,t}$$

where all variables are defined as in Equation (1) and $Interaction\_1_i$ is a dummy equal to 1 if the target falls in the leisure category and zero otherwise, $Interaction\_2_i$ is a dummy equal to 1 if the target falls in the durable category and zero otherwise and $Interaction\_3_i$ is a dummy equal to 1 if the target falls in the "other" category and zero otherwise. By excluding the generic saving category *savings*, we take it as the base case.

The first column of Table 6 shows economically small effects related to the type of target. All interactions are negative but only the ones for "leisure" and "other" are statistically significant. Economically the effects are negative but not very large, suggesting that the specific purpose of the target does not have an effect on users' saving behavior, consistent with the survey results and in contrast with the predictions of a mental accounting channel.

The second column focuses on the horizon of the goal and $Interaction\_1_i$ identifies targets with horizons of less than (or equal to) a year. We therefore take targets with horizons greater than one year as the base case. Individuals save significantly more when they set targets of short horizon: the coefficient on the short horizon interaction is positive and statistically significant, but the effect is economically small. Combined with the survey findings, these results suggest that the shorter horizon makes it easier for individuals to monitor their progress.

The third column focuses on the amount of the goal and $Interaction\_1_i$ identifies targets smaller than €1,000. We therefore take targets with horizons greater than €1,000 as the base case. We do not find any relation between the size of the target and the saving behavior.

5.3 Likelihood of Achieving Goals

We start by presenting basic facts on goal attainment. For this analysis, we restrict the sample to the targets that had a deadline before the date on which we obtained the data, November 2nd, 2019. We have a total of 14,521 goals. Of these 4,088 (28.15%) were attained within the deadline, 3,568 (24.57%) were closed and unattained, and 6,865 (47.28%) were overdue but still active. The closed and overdue goals are generally rather far from being attained. On average, only 11% of the target was funded for closed targets. The corresponding percentage for overdue targets was 13.1%.

To assess whether the specificity of the goal relates to the probability of achieving it, we estimate a logit model of the following form:

$$y_i = \begin{cases} 1 & \text{if goal } i \text{ is achieved} \\ 0 & \text{otherwise,} \end{cases}$$

and

$$p_i = Pr\left(y_i = 1 | \boldsymbol{x}_i\right) = \frac{\exp\left(\boldsymbol{x}_i'\boldsymbol{\beta}\right)}{1 + \exp\left(\boldsymbol{x}_i'\boldsymbol{\beta}\right)}, \tag{6}$$

The vector $\mathbf{x}_i$ contains dummies for the following goals: hobby, travel, other, car, house, and generic savings (the base case). Other control variables included are: the 1) the log transformation of the goal horizon—measured in days; 2) the log amount of the goal; 3) the age of the person setting the goal; and 4) whether a person was a male. All continuous variables are standardized so they have a standard deviation of one.

The results are reported in Table 7. The first and second columns report coefficients for logit models, while columns 3 and 4 provide results for linear probability models. The second and fourth columns include monthly time effects and province fixed-effects. Across all specifications, we find virtually identical results. Starting from goal categories, specific goals are negatively related to the probability of success. The hobby, travel, other, car, and house categories have a probability of success 10 percentage points lower than the generic saving (the base category). As expected, the control variables size and horizon have a large impact on the probability of success. A one standard deviation increase in the size of the goal decreases the probability of success by 13 percentage points, and a one standard deviation increase in the horizon decreases the probability of success by 3 percentage points. This implies that short-term and small goals are the most likely to be achieved. Also, older users are more likely to achieve their goals, possibly because they are less financially constrained, more experienced in designing realistic goals, and/or more responsible with their finances. Finally, men are significantly less likely to achieve their goals, possibly because they are overconfident and less responsible than females.

5.3.1 Alternative Goal Categorization.

The pre-specified categories available on the App might not fully span the purpose of the goals set by users. Moreover, after describing the goal, users might be careless in selecting the category to which it belongs to (i.e. she may select the category "savings" even if the name of the target is "Trip to New York"). In both cases, the goal indicator variables in Equation 6 would contain measurement error which, if large enough, could bias the coefficient estimates. To address these concerns, in Online Appendix A we use data from the text entered by the users to describe the goals and re-categorize the goals both manually and using automated machine learning methods. We find that these alternative goal categorizations return the same results.

5.3.2 Additional Evidence Provided by the Survey.

Because observational data do not allow us to fully infer what factors determine users' likelihood of achieving their saving goals, we complement these results with survey data. In panel b (c) of Figure 8 we ask users which factors they believe have determined the success (failure) of their goals. Again, consistent with Table 7, when asked about their attained goals, only 18% of the users claim the purpose of the goal was the most decisive factor. About 54% of respondents claim instead that fund returns and their ability to monitor the progress towards their goal have played a more important role. When asked about the goals they have not attained, a considerable fraction of respondents (27.1%) are unable to provide a precise answer and choose the "Do not know" option. Of the remaining options, the most common is "unexpected expenses" (30%) followed by a too ambitious goal (15.8%), poor fund returns (9%), and lost motivation (7%).

5.3.3 Dynamics of Savings towards the Goal.

The results on goal achievement reported above are silent about the time-dynamics of saving behavior. For example, do individuals who do not achieve goals get fatigued over time and stop saving as time passes by or do they fail to be on track from the very beginning? For those who achieve their goals, do they save more than they should at the beginning and lose steam as time goes by? Providing answers to these questions is crucial to help individuals design achievable and motivating saving goals.

To this end, we first restrict the sample to the 14,521 goals goals with deadline before the date on which we obtained the data (November 2nd, 2019). To ease the comparison across goals with different characteristics, we normalize to one both the horizon and the amount of each goal. For each goal, we then cumulate the amount deposited by the user over time so that we have the percentage of the goal completed as a function of the percentage of the time elapsed. Finally, we compute semi-parametric estimates of the average percentage completion across all goals, as a function of the time elapsed.

Results are reported in Figure 9. The first subfigure reports the results pooling across all goals (in black) and separating goals that are attained (in blue), and goals that are unattained (in green).[13] While difficult to discern, the 95% confidence intervals are also reported with dashed lines. We highlight several facts. First, pooling across goals, individuals save only 20% of the goal amount by the time the goal expires, on average. Second, unattained goals are very different from attained ones right from their inception. In the first, individuals under-save from the very beginning—the green line is always below the red 45 degree line—while, if anything, individuals over-save in the first part of the goals that are attained. Individuals over-save up to the 40-th percentile of the time duration of the goal and progressively decrease their saving rate thereafter, only to increase it at the very end to complete the goal.

Overall, the results in this section suggest that the type of goal set by the individuals is barely related to the probability of achieving the target, and that the stronger savings in connection with the adoption of targets documented in the previous section is likely due to the fact that the App allows individuals to monitor and track their progress towards their savings goals.

[13]Recall that approximately 24.57% of these goals were attained within the deadline.

6 Understanding the Economic Effects of Goal Setting on Saving

In this section we undertake two final exercises to understand whether goal-setting can help those individuals the literature has found to be particularly at risk of under-saving (Section 6.1) and whether the increases in savings associated with goal-setting are economically large relative to users' income (Section 6.2).

6.1 Goal-Setting and Other Determinants of Under-saving

The literature has identified a number of reasons that explain individuals' under-saving. These explanations range from low levels of financial literacy (Lusardi, 2008 and Lusardi and Mitchell, 2014), high impatience (Laibson, 1997), low levels of overall education (Cole, Paulson, and Shastry, 2014 and Bernheim, Garrett, and Maki, 2001), low income and socioeconomic status (Banerjee and Mullainathan, 2010, Duflo and Banerjee, 2011 and Kuhnen and Miu, 2017), and user's inattention (Karlan et al., 2016b and Gargano and Rossi, 2018). In this section, we estimate whether goal-setting can help the users the literature has identified to have the lowest propensity to save.

To this end, in Table 8 we estimate:

$$M\_Net\_Deposits_{i,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,t} + \gamma\ Target\_Dummy_{i,t} \times Inter_i + \epsilon_{i,t},$$

where all regressors are defined as in Equation (1) but now the dummy $Inter_i$ is equal to one for those individuals with high levels of financial literacy (column 2), high patience (column 3), high education (column 4), high income (column 5) and high degree of attention on the App before engaging in goal-setting (column 6). We classify as highly financially literate those individuals who claim to have either medium-high and high knowledge of investing and answer correctly the Lusardi-Mitchell question in the Survey—see Figure A.2 for details. High-patience individuals are defined as those who answer €101 or €103 to the patience question, high-education individuals are those who completed at least a bachelor degree as their highest level of education and high income individuals are those who earn more than €2,000 in net income per month—see Figure A.2 for details. Finally, the sixth column identifies as high attention those users who, before engaging in goal-setting, are above the median in terms of number of logins.

The coefficient of interest in all our specifications is β, which estimates the effect of goal-setting on those individuals who are at the highest risk of under-saving, according to the literature. The γ coefficient instead estimates the differential effect of goal-setting on those individuals that the literature has identified as high savers compared to low-savers.

Given that five of the six conditioning variables are available only for a subset of users who complete the survey, the first column does not include any interaction and simply re-estimates the specification in column 1 of Table 2 on the sample of survey respondents. It establishes that the results on this subsample are very similar to the ones computed on the full sample of App users, allaying concerns of sample selection biases in the population of survey respondents.

Moving to the remaining specifications, across columns 2 through 6 the β coefficients show that even those who are more prone to under-saving increase significantly their saving after engaging in goal-setting. If anything, in the majority of instances, we find that goal setting has the same effect across all users. This can be seen in columns 3 through 5, where the γ coefficient is statistically insignificant, suggesting there is no differential increase in saving associated with goal setting for the individuals with high patience, high education and high income. In column 6 instead the results suggest that high-attention individuals benefit less from goal-setting, consistent with the monitoring mechanism we have highlighted in Section 5: the users who were already monitoring their saving before engaging in goal-setting are the ones who experience less of a boost in saving after setting goals. The only conditioning variable for which we find a positive γ coefficient is financial literacy. The interaction between financial literacy and goal setting is estimated to increase the effect of goal-setting by 19.82/31.94=62%, which is statistically and economically large and suggests FinTech apps could be deployed in conjunction with financial literacy to amplify their effectiveness.

Overall, the results in this section suggest that goal-setting can be used to help the portion of the population that is at highest risk of under-saving.

6.2 The Economic Benefits of Goal-Setting Relative to Users' Income

In this last section, we quantify the economic effects of goal-setting in boosting savings for the subset of App users who participate in the survey and provide monthly income information. This allows us to compute the saving rate for each individual as the percentage of her net income.[14] We construct, first, individuals' saving rate before adopting goal setting. Second, the implied saving rate associated with saving goals: the saving rate individuals would like to have as they are planning for the future. Third, the saving rate users achieve when they use goal-setting, which is the result of planning and having the freedom to deviate from the original plan.

The difference between the saving rate implied by goals and the saving rate when not setting goals can be thought to represent the potential effect of goal setting in increasing savings rates. The difference between the saving rate implied by the goals and the saving rate obtained with goal-setting can instead be taken to represent the reduction in saving (leakage) individuals experience with soft commitments. Finally, the difference between the saving rate with goals and the saving rate without goals measures the realized effect of goal setting. We display schematically these relations in the diagram reported below:

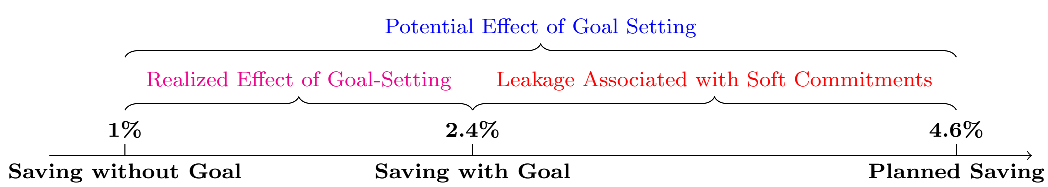

The median individual saves 1% of her net income when she does not set goals. The planned percentage savings associated with goals equal 4.6% of her net income. Finally, the percentage savings realized when setting goals equal 2.4% of her income. We obtain virtually identical values when we restrict the analysis to only those individuals who do not save anything outside the App: 1%, 6%, and 2%, respectively.[15] These differences in savings rate highlight two facts. First, there is substantial

[14] Note that the survey asks users for their net income in discrete intervals—see Figure A.2 for details. We take the mid-point of each interval and assign €4,500 to those who respond they earn more than €4,000.

[15] We work with medians, because they are less affected by outliers. If we use averages we obtain that the average saving rates as percentage of income are 2.1% for individuals who do not set goals and 4.0% for individuals who save with goals—in line with the median values reported in the main body of the paper. The savings rate implied by savings goal plans is instead higher and averages 11.7%, mainly because some individuals set unrealistically high saving goals and the large positive skewness increases the averages.

leakage associated with the soft commitments individuals engage in. Second, goal setting is rather helpful in helping individuals increase their baseline saving rate.

Economically, the effects are rather large. Take an individual who invests in an all equity mutual fund and therefore can expect a 6% annual risk premium. This individual would save one month worth of her net income in 55 months (almost five years). Under the saving rate implied by the saving plan (4.6%) the very same individual would achieve the same overall saving in 19 months (a little less than two years). Finally, if she was engaging in goal-setting, she would achieve that overall saving in 33 months (a little less than 3 years).

7 Conclusions

We study the effectiveness of *soft*, *self-designed* commitment devices in saving decisions using data from an App developer that provides its users with simple budgeting solutions in the form of saving goals. We complement the administrative App data with a survey administered on a sample of App users to uncover the economic mechanism underlying individuals' goal-setting choices.

We establish that setting goals increases individuals' saving rate and show that the effect is causal using a difference-in-differences identification strategy that exploits the random assignment of users into a group of beta-testers who can set goals and a group of users who cannot. We also show that the increased savings within the App do not come at the expense of reduced savings outside the App and that goal setting is very effective for the portion of the population at greater risk of under-saving.

We explore the economic channels of our results by matching App user survey responses to their behavior and highlight the importance of a monitoring channel, consistent with models where agents experience disutility from falling short of their goal and goal setting reducing individuals' limited attention.

Taken together, our findings indicate that soft, self-designed commitment devices have an overall positive effect on saving. This is reassuring, given that a large number of FinTech firms and large financial institutions have been introducing these tools in recent years. At the same time, individuals oftentimes fail to reach their goals, suggesting there is considerable scope to increase the effectiveness of goal-setting by helping individuals set well-calibrated goals, possibly using robo-advising tools

(D'Acunto and Rossi, 2021). In addition, given that mobile phone usage is becoming widespread across developing countries—see Jack and Suri (2014) and Duflo and Banerjee (2011) for examples—the soft, self-designed commitment devices we study in this paper could be deployed to help low income individuals, who are arguably the ones who need these FinTech tools the most.

References

Agarwal, S., W. Qian, Y. Ren, S.-T. Tsai, and B. Y. Yeung. 2020. The real impact of fintech: Evidence from mobile payment technology. Working Paper.

Agarwal, S., B. Y. Yeung, and X. Zou. 2019. Mobile wallet and entrepreneurial growth. *AEA Papers and Proceedings* 109:48–53.

Ashraf, N., D. Karlan, and W. Yin. 2006. Tying odysseus to the mast: Evidence from a commitment saving product in the philippines. *Quarterly Journal of Economics* 121:635–72.

Banerjee, A., and S. Mullainathan. 2010. The shape of temptation: Implications for the economic lives of the poor. NBER Working Paper.

Banks, J., R. Blundell, and S. Tanner. 1998. Is there a retirement-savings puzzle? *American Economic Review* 4:769–88.

Bénabou, R., and J. Tirole. 2004. Willpower and personal rules. *Journal of Political Economy* 112:848–86.

Benartzi, S., and R. H. Thaler. 2007. Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives* 21:81–104.

Bernheim, B. D., D. M. Garrett, and D. M. Maki. 2001. Education and saving:: The long-term effects of high school financial curriculum mandates. *Journal of Public Economics* 80:435–65.

Bernheim, B. D., J. Skinner, and S. Weinberg. 2001. What accounts for the variation in retirement wealth among us households? *American Economic Review* 91:832–57.

Beshaers, J., J. J. Choi, C. Harris, D. Laibson, B. C. Madrian, and J. Sakong. 2020. Which early withdrawal penalty attracts the most deposits to a commitment savings account? *Journal of Public Economics* 183:1041–4.

Beshears, J., J. J. Choi, D. Laibson, B. C. Madrian, and W. L. Skimmyhorn. Forthcoming. Borrowing to save? the impact of automatic enrollment on debt. *Journal of Finance* .

Beshears, J., K. L. Milkman, and J. Schwartzstein. 2016. Beyond beta-delta: The emerging economics of personal plans. *American Economic Review* 106:430–34.

Brauer, K., A. Hackethal, and T. Hanspal. 2020. Consuming dividends. Working Paper.

Breza, E., and A. G. Chandrasehkar. 2019. Social networks, reupation and commitment: Evidence from a savings monitors experiment. *Econometrica* 87:175–2016.

Brown, J. R., and A. Previtero. 2020. Saving for retirement, annuities, and procrastination. Working Paper.

Bryan, G., D. Karlan, and S. Nelson. 2010. Commitment devices. *Annu. Rev. Econ.* 2:671–98.

Carlin, B., A. Olafsson, and M. Pagel. 2019. Fintech and financial fitness in the information age. Working Paper.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick. 2009. Reinforcement learning and savings behavior. *The Journal of finance* 64:2515–34.

Choi, J. J., and A. Z. Robertson. 2020. What matters to individual investors? evidence from the horse's mouth. *The Journal of Finance* 75:1965–2020.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. Weber. 2020. Forward guidance and household expectations. Tech. rep., National Bureau of Economic Research.

Coibion, O., Y. Gorodnichenko, and M. Weber. 2019. Monetary policy communications and their effects on household inflation expectations. Tech. rep., National Bureau of Economic Research.

Cole, S., A. Paulson, and G. K. Shastry. 2014. Smart money? The effect of education on finanical outcomes. *Review of Financial Studies* 27:2022–51.

D'Acunto, F., U. Malmendier, and M. Weber. Forthcoming. Gender roles and the gender expectations gap. *Proceedings of the National Academy of Sciences* .

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

D'Acunto, F., T. Rauter, C. Scheuch, and M. Weber. 2019. Perceived precautionary savings motives: Evidence from FinTech. Working Paper.

D'Acunto, F., and A. G. Rossi. 2021. New frontiers of robo-advising: Consumption, saving, debt management, and taxes. *Machine Learning in Financial Markets: A Guide to Contemporary Practice, Cambridge University Press,* .

———. Forthcoming. Robo-advising. *Palgrave Macmillan Handbook of Technological Finance* .

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

Das, S. R. 2019. The future of FinTech. Working Paper.

DellaVigna, S., and U. Malmendier. 2006. Paying not to go to the gym. *american economic Review* 96:694–719.

Di Maggio, M., and V. Yao. 2020. Fintech borrowers: Lax-screening or cream-skimming? Working Paper.

Duflo, E., and A. Banerjee. 2011. *Poor economics*, vol. 619. PublicAffairs.

Duflo, E., M. Kremer, and J. Robinson. 2011. Nudging farmers to use fertilizer: Theory and experimental evidence from kenya. *American Economic Revirw* 6:2350–90.

Dupas, P., and J. Robinson. 2013. Why don't the poor save more? Evidence from health savings experiments. *American Economic Review* 103:1138–71.

Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde. 2018. Global evidence on economic preferences. *The Quarterly Journal of Economics* 133:1645–92.

Fuster, A., P. Godsmith-Pinkham, T. Ramadorai, and A. Walther. 2018. Predictably unequal? The effects of machine learning on credit markets. Working Paper.

Gargano, A., and M. Giacoletti. 2021. Cooling auction fever: Underquoting laws in the housing market. Working Paper.

Gargano, A., M. Giacoletti, and E. Jarnecic. 2019. Local experiences, attention and spillovers in the housing market. Working Paper.

Gargano, A., and A. G. Rossi. 2018. Does it pay to pay attention? *The Review of Financial Studies* 31:4595–649.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus. Forthcoming. Five facts about beliefs and portfolios. *American Economic Review* .

———. 2021. The joint dynamics of investor beliefs and trading during the COVID-19 crash. *Proceedings of the National Academy of Sciences* 118:e2010316118–. doi:10.1073/pnas.2010316118.

Giné, X., D. Karlan, and J. . Zinman. 2010. Put your money where your butt is: a commitment contract for smoking cessation. *American Economic Journal: Appliled Economics* 2:213–35.

Guiso, L., P. Sapienza, and L. Zingales. 2008. Trusting the stock market. *Journal of Finance* 63:2557–600.

Hsiaw, A. 2013. Goal-setting and self-control. *Journal of Economic Theory* 148:601–26.

Jack, W., and T. Suri. 2014. Risk sharing and transactions costs: Evidence from kenya's mobile money revolution. *American Economic Review* 104:183–223.

John, A. 2020. When commitment fails: evidence from a field experiment. *Management Science* 66:503–29.

Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman. 2016a. Getting to the top of mind: How reminders increase saving. *Management Science* 62:3393–411.

Karlan, D., M. McConnell, M. S, and J. Zinman. 2016b. Getting to the top of mind: How reminders increase saving. *Management Science* 62:3393–672.

Kuchler, T., and M. Pagel. 2021. Sticking to your plan: The role of present bias for credit card paydown. *Journal of Financial Economics* 139:359–88.

Kuhnen, C. M., and B. Knutson. 2011. The influence of affect on beliefs, preferences, and financial decisions. *Journal of Financial and Quantitative Analysis* 46:605–26.

Kuhnen, C. M., and B. T. Melzer. 2018. Noncognitive abilities and financial delinquency: The role of self-efficacy in avoiding financial distress. *The Journal of Finance* 73:2837–69.

Kuhnen, C. M., and A. C. Miu. 2017. Socioeconomic status and learning from financial information. *Journal of Financial Economics* 124:349–72.

Laibson, D. 1997. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics* 112:443–78.

Laudenbach, C., J. Pirschel, and S. Siegel. 2018. Personal communication in a FinTech world: Personal communication in a fintech world: Evidence from loan payments. Working Paper.

Liu, H., C. Peng, W. A. Xiong, and W. Xiong. Forthcoming. Taming the bias zoo. *Journal of Financial Economics* .

Locke, E. A., and G. P. Latham. 1990. *A theory of goal setting and task performance.* NJ: Prentice Hall.

———. 2002. Building a practically useful theory of goal setting and task motivation. *American Psychologist* 57:705–17.

———. 2006. New directions in goal-setting new directions in goal-setting theory. *Current Directions in Psychological Science* 15:265–8.

Loh, R. K., and H.-S. Choi. 2020. Physical frictions and digital banking adoption. Working Paper.

Lusardi, A. 2008. Household saving behavior: The role of financial literacy, information, and financial education programs. NBER Working Paper.

Lusardi, A., and O. S. Mitchell. 2007. Baby Boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics* 54:205–24.

———. 2017. How ordinary consumers make complex economic decisions: financial literacy and retirement readiness. *Quarterly Journal of Finance* 7:1750008–.

———. 2014. The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature* 52:5–44.

Madrian, B. C., and D. F. Shea. 2001. The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics* 116:1149–87.

Medina, P. C., and M. Pagel. 2021a. Does saving cause borrowing? *National Bureau of Economic Research* .

Medina, P., and M. Pagel. 2021b. Does saving cause borrowing? *Working Paper* .

O'Donoghue, T., and M. Rabin. 1999. Doing it now or later. *American economic review* 89:103–24.

Rossi, A. G., and S. Utkus. 2020. Who benefits from robo-advising? Evidence from machine learning. Working Paper.

Shefrin, H. M., and R. H. Thaler. 1988. The behavioral life-cycle hypothesis. *Economic inquiry* 26:609–43.

Siegel, J. 1998. *Stocks for the long run.* New York: McGraw-Hill.

Thaler, R. 1985. Mental accounting and consumer choice. *Marketing science* 4:199–214.

Thaler, R. H. 1990. Anomalies: Saving, fungibility, and mental accounts. *Journal of economic perspectives* 4:193–205.

———. 1999. Mental accounting matters. *Journal of Behavioral decision making* 12:183–206.

Thaler, R. H., and S. Benartzi. 2004. Save more tomorrow$^{\text{TM}}$: Using behavioral economics to increase employee saving. *Journal of political Economy* 112:S164–87.

Thaler, R. H., and H. M. Shefrin. 1981. An economic theory of self-control. *Journal of political Economy* 89:392–406.

Thaler, R. H., and C. R. Sunstein. 2009. *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.

Figure 1: This Figure displays the four steps required to create a saving goal: (a) naming the goal; (b) choosing between six categories: Hobby (Tempo Libero), Travel (Viaggi), Other (Altro), Car (Veicoli), House (Casa) and Savings (Risparmio); (c) choosing a target amount (e.g. € 3000) and a deadline (e.g. 30/08/2019); (d) Linking the goal to an image.

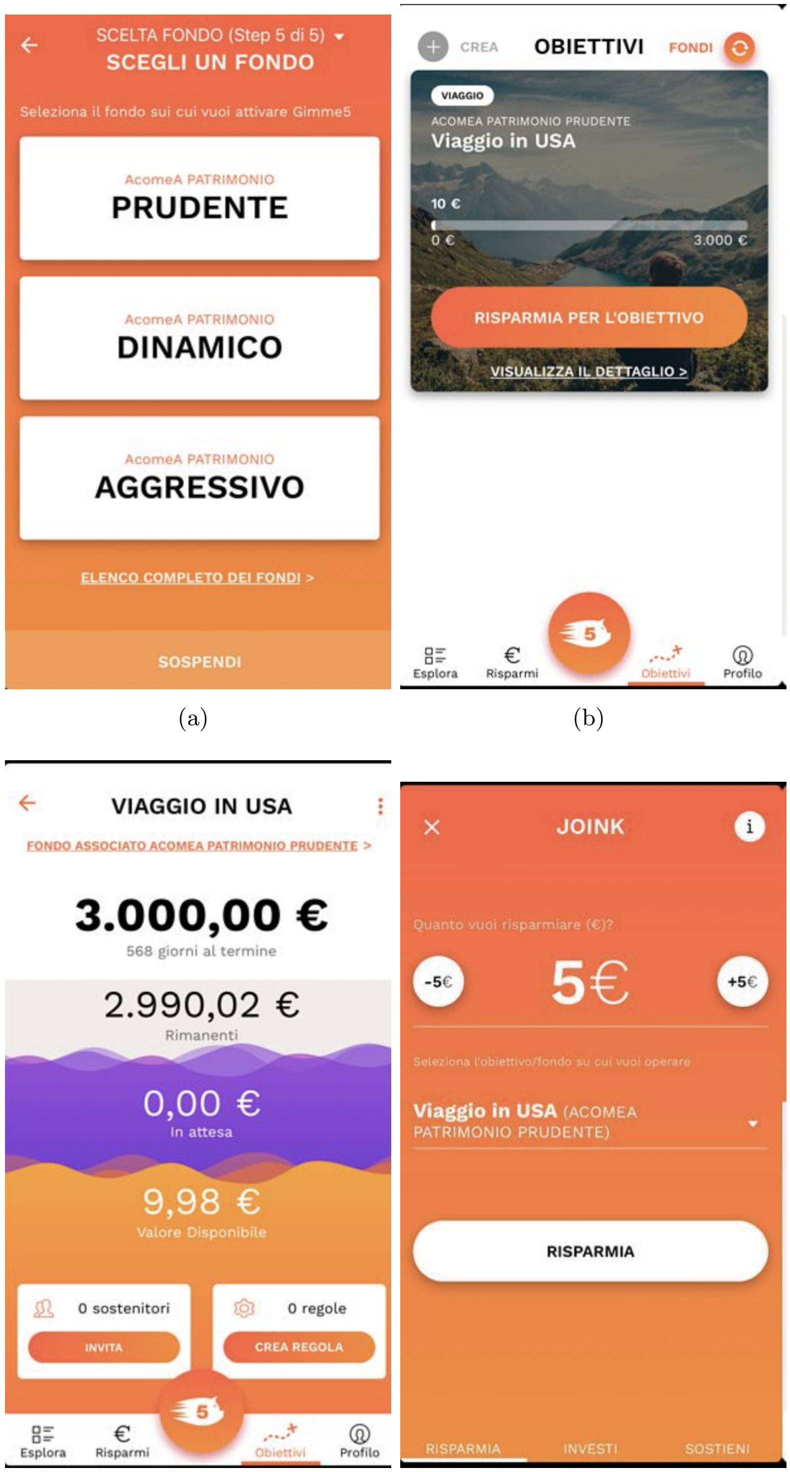

(a) (b)

(c) (d)

Figure 2: This Figure displays the section of the App where users choose the mutual fund for their savings (a); monitor the goal progression (b and c); pledge money with the "Joink" function (d).

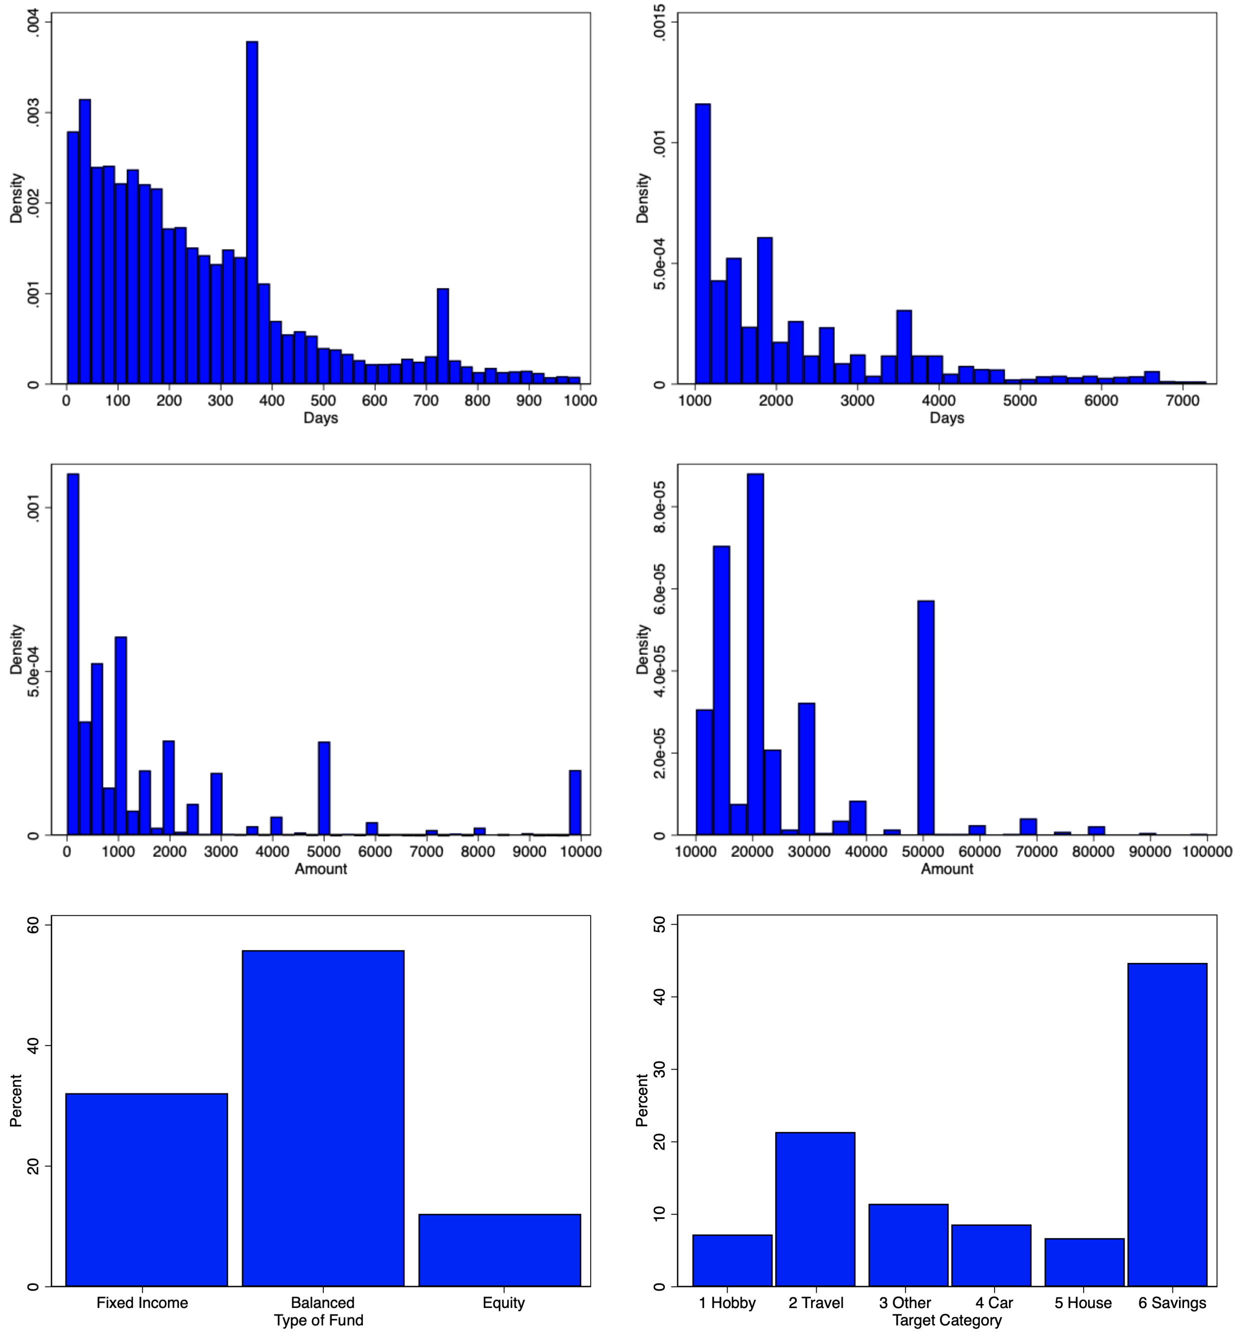

Figure 3: This figure displays the cross-sectional distribution of four key features of the saving goals set by users. The upper subfigures report the horizon of each goal, computed as the difference (in days) between the targeted day and the day when the goal was set. Because the distribution displays a thin right tail, we display it in two parts: the upper-left subfigure is associated with horizons shorter than 1,000 days while the upper-right one is associated with horizons longer than 1,000 days. The middle subfigures report the amounts (in €) users set to save. The left subfigure refers to goals less than or equal to €10,000, while the right subfigure refers to goals grater than €10,000. The bottom-left subfigure displays the type of funds users invest their savings in: Fixed Income, Balanced and Equity funds. The bottom-right subfigure displays the categories of goals users can choose from: Hobby, Travel, Other, Car, House and Savings.

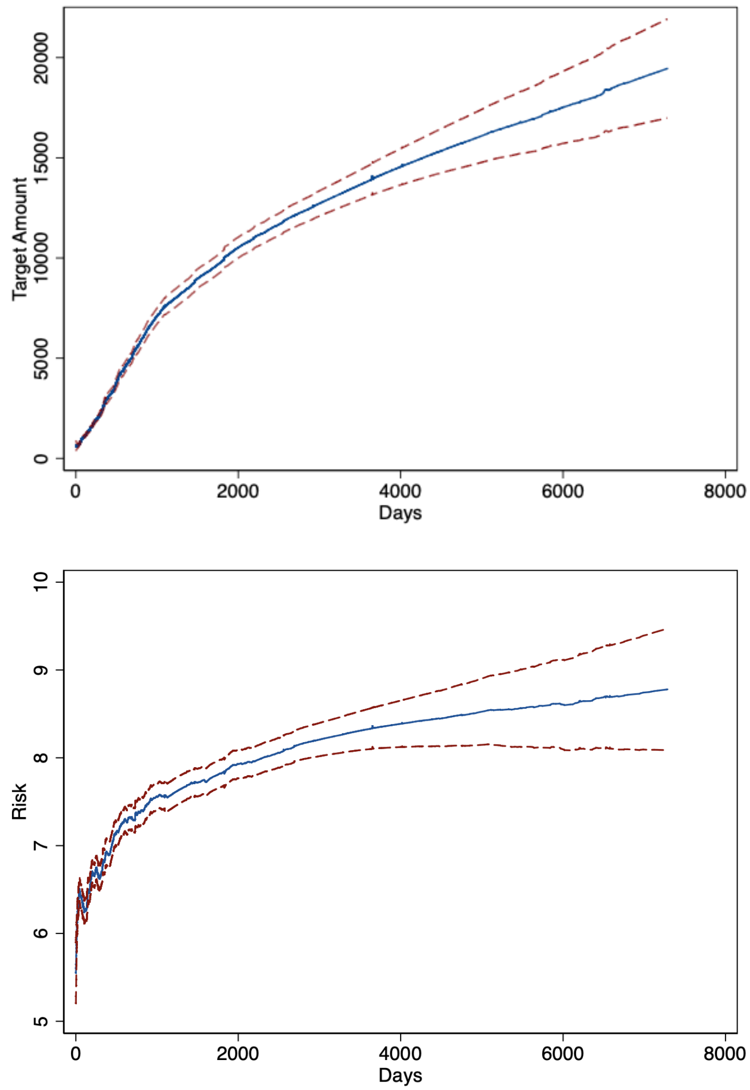

Figure 4: This figure displays the relation between the goal amount (upper subfigure), risk (bottom subfigure), and horizon. The blue lines denote the nearest neighbor estimate of the conditional mean, while the red dashed lines denote the 95% confidence intervals. Risk is measured as the annualized volatility of the fund.

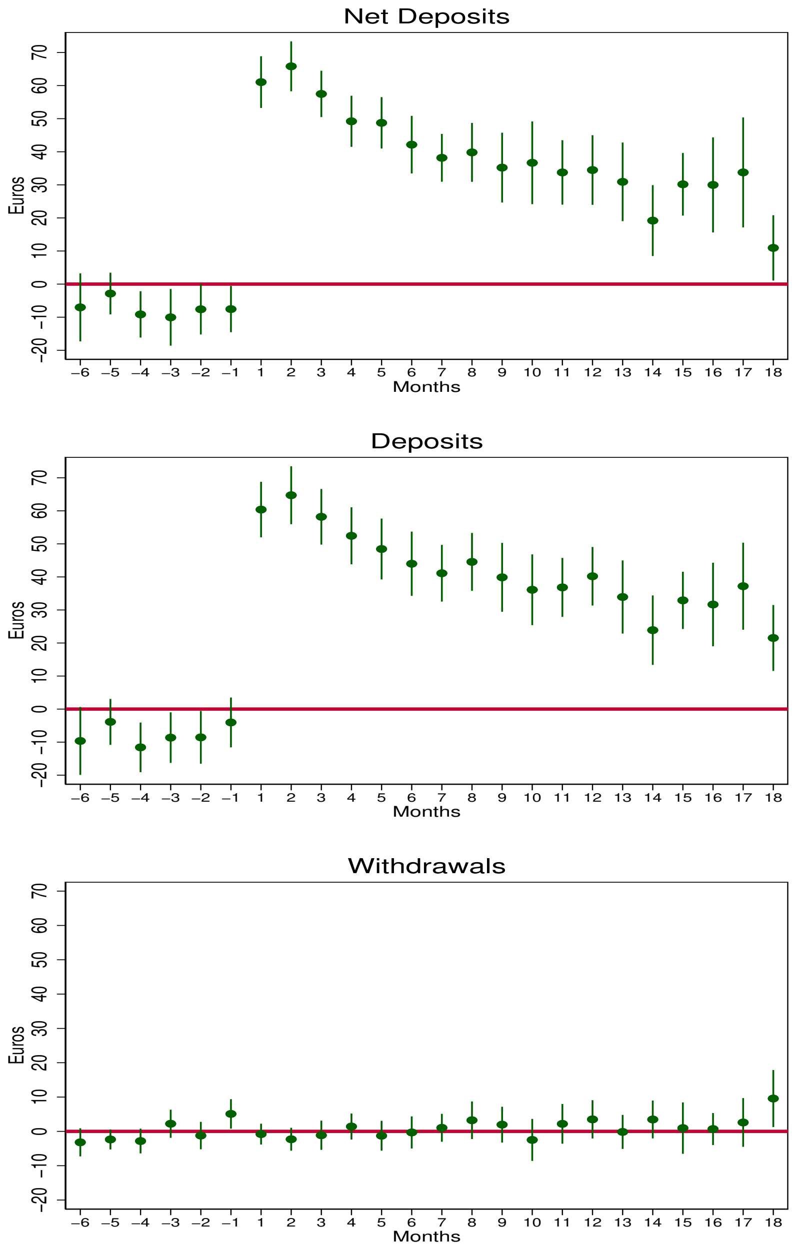

Figure 5: This figure reports coefficient estimates and 95% confidence interval of the β_s coefficients from the following baseline panel regression:

$$M\_Net\_Deposits_{i,j,t} = \alpha_{i,t} + \sum_{s=-6}^{18} \beta_s \, Target\_Dummy_{i,j,s,t} + \epsilon_{i,j,t}.$$

where $M\_Net\_Deposits_{i,j,t}$ denotes the monthly amount (in €) transacted by user i in account j at time t; $\alpha_{i,t}$ identify user-time fixed-effects; $Target\_Dummy_{i,j,s,t}$ is a dummy variable equal to 1 if account j associated with user i is using a target for the $s-th$ month in the calendar month t. The result in the upper subfigure are based on $Net\ Deposits$, i.e. the difference between total monthly $Deposits$ and $Withdrawals$, while the results in the second and bottom subfigures only use $Deposits$ and $Withdrawals$, respectively. The standard errors are double-clustered at the user and month levels.

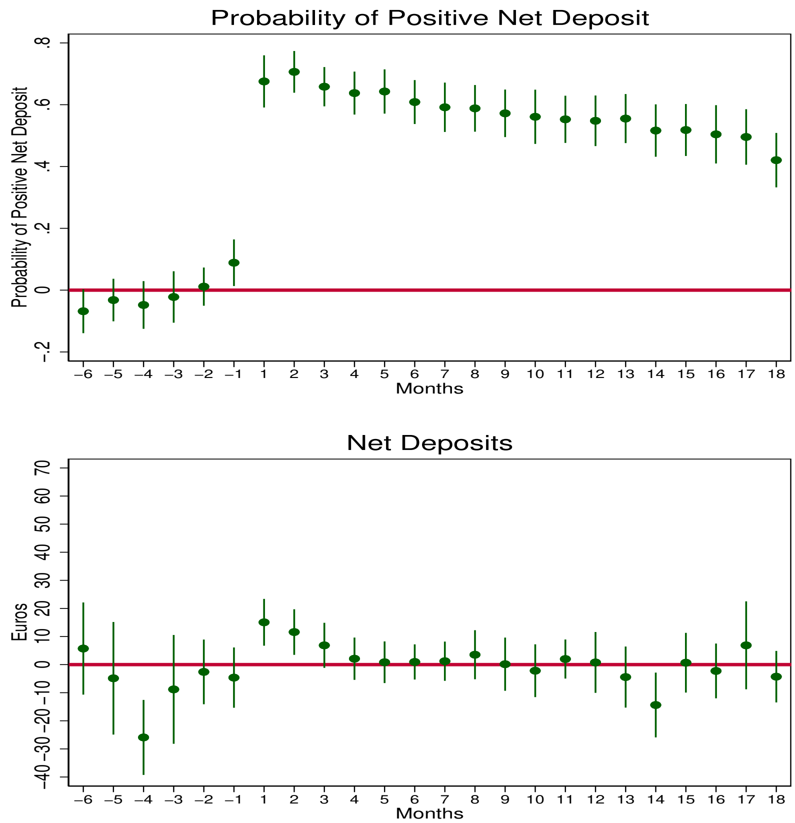

Figure 6: This figure reports coefficient estimates and 95% confidence interval of the β_s coefficients from the following regressions:

$$IND_{i,j,t} = \alpha_{i,t} + \sum_{s=-6}^{18} \beta_s \; Target\_Dummy_{i,j,s,t} + \epsilon_{i,j,t} \qquad (Top\ Subfigure)$$

$$M\_Net\_Deposits^{+}_{i,j,t} = \alpha_{i,t} + \sum_{s=-6}^{18} \beta_s \; Target\_Dummy_{i,j,s,t} + \epsilon_{i,j,t} \qquad (Bottom\ Subfigure)$$

where $IND_{i,j,t}$ is an indicator variable equal to 1 if user i deposited a positive amount in account j in month t and 0 otherwise and $M\_Net\_Deposits^{+}_{i,j,t}$ denotes the monthly net deposit (in €) by user i in account j at time t, conditional on $IND_{i,j,t}$ being equal to 1; $\alpha_{i,t}$ identify user-time fixed-effects; $Target\_Dummy_{i,j,s,t}$ is a dummy variable equal to 1 if account j associated with user i is using a target for the $s-th$ month in the calendar month t. The standard errors are double-clustered at the user and month levels.

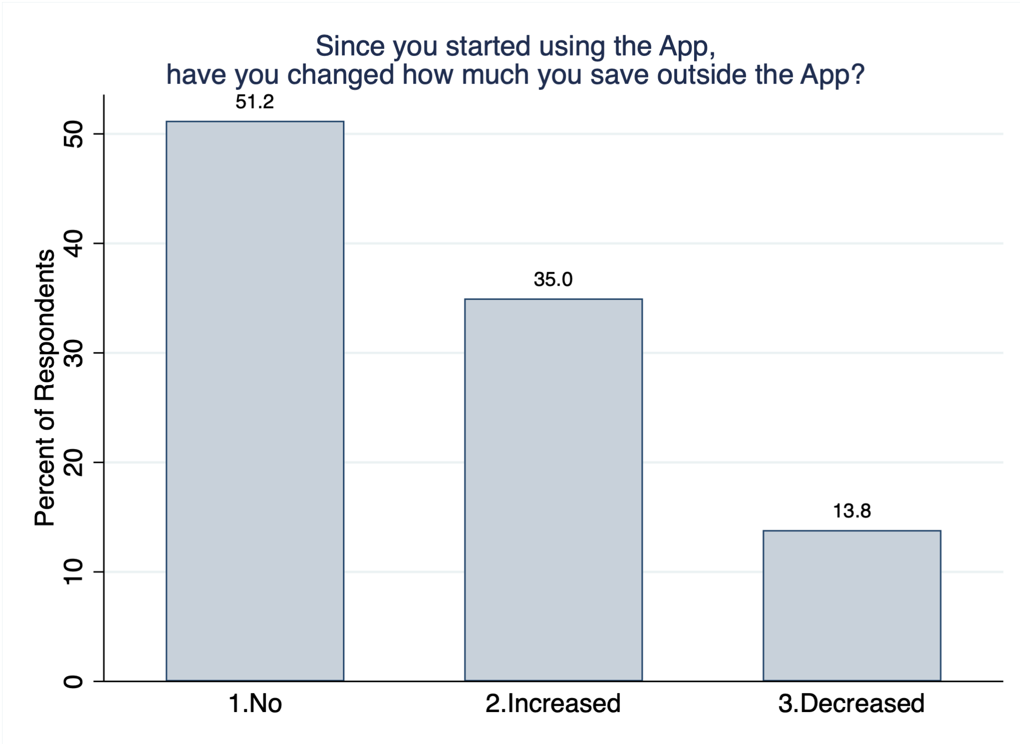

(a)

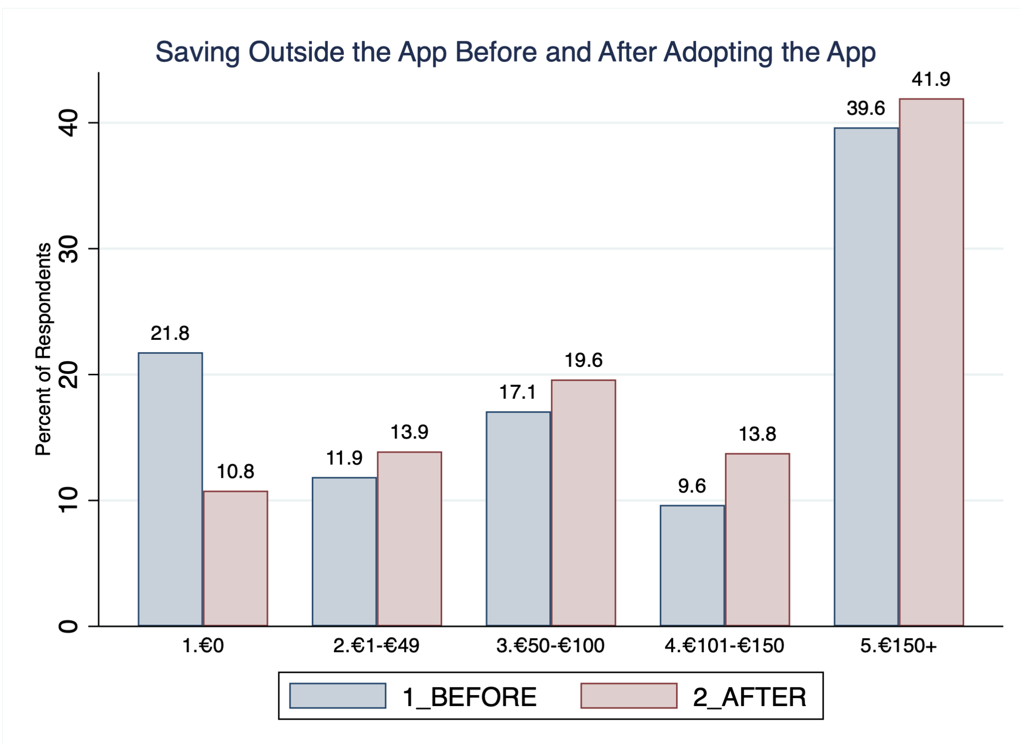

(b)

Figure 7: This Figure shows the distribution of the answers to the survey questions relative to users' saving behavior outside the App before and after its adoption. The question associated with Panel a) is *"Since using the App, have you changed how much you saved outside the App?"*. The questions associated with Panel b) are *"Before using the App, how much have you saved on average per month outside the App?"* (blue bars) and *"Since using the App, how much have you saved on average per month outside the App?"* (red bars). We display the percentage of respondents at the top of each bar.

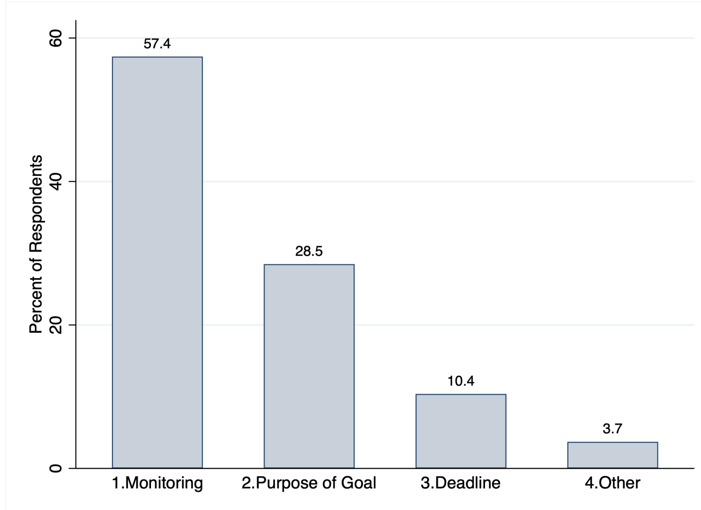

(a) Motivating Aspects of Saving Goals

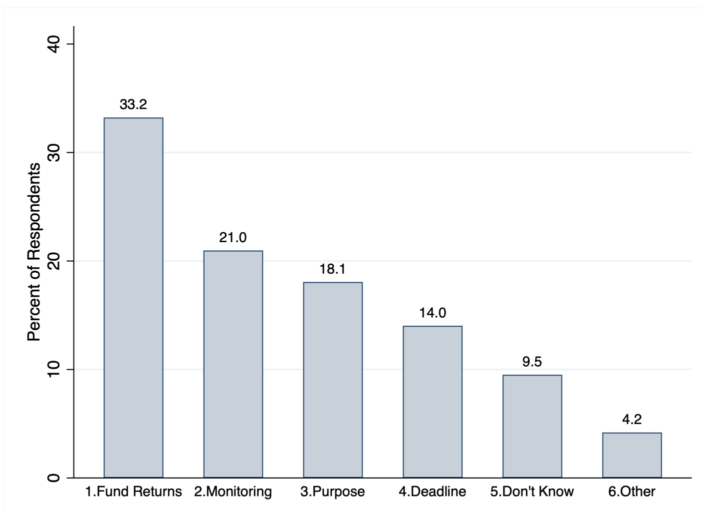

(b) Attained Goals

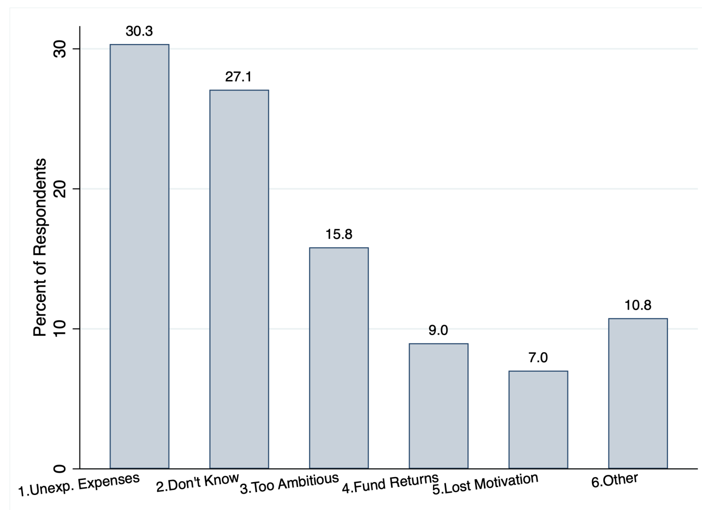

(c) Not Attained Goals

Figure 8: This Figure shows the distribution of the answers to the survey questions relative to the use of saving goals. The question associated with Panel a) is *"In your experience, what is the most motivating aspect of having a saving goal?"*. The question associated with Panel b) is *"Think about the saving goals you have achieved. In your opinion, what has been the most decisive factor?"*. The question associated with Panel c) is *"Think about the saving goals you have NOT achieved. In your opinion, what has been the most decisive factor?"*. We display the percentage of respondents at the top of each bar.

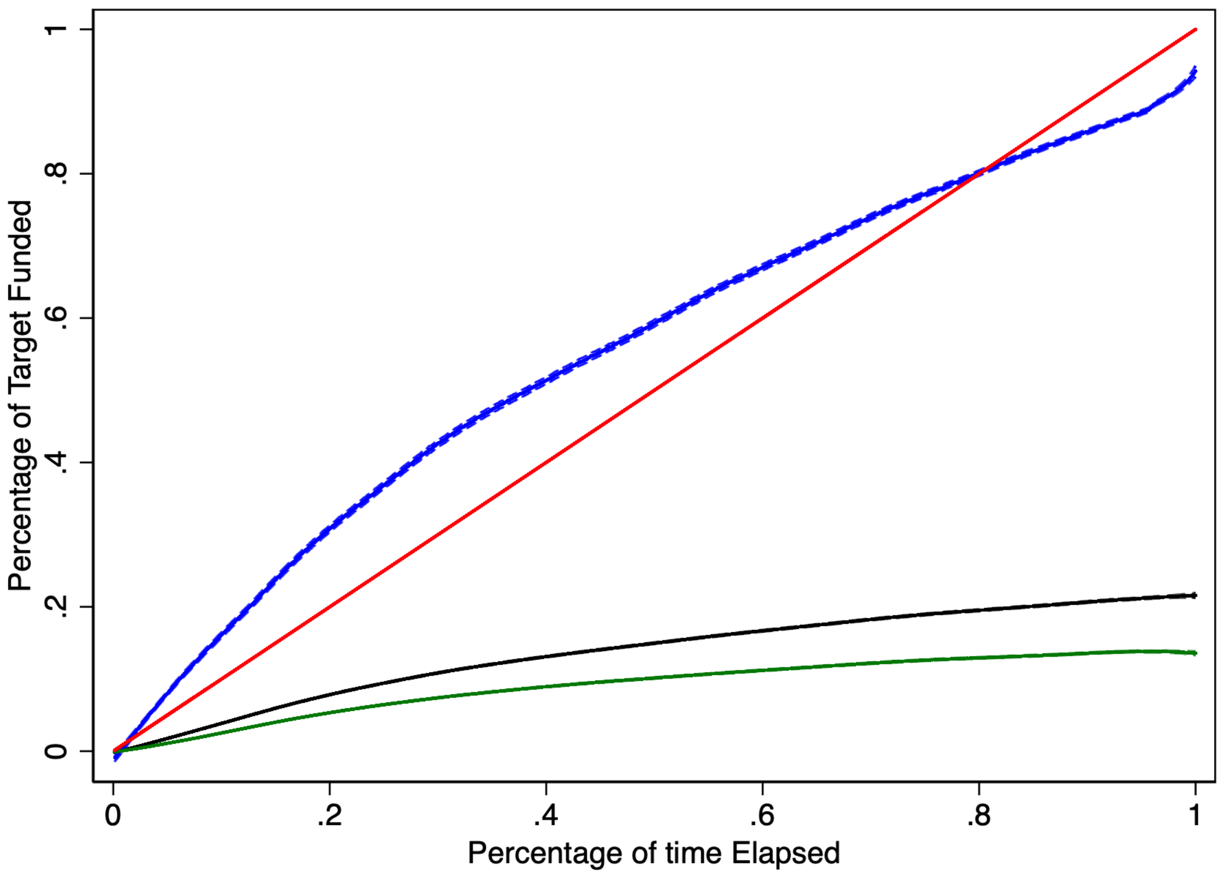

Figure 9: This figure analyzes the dynamics of saving when individuals set goals. The sample consists of the 14,521 goals with deadline before the date on which we obtained the data (November 2nd, 2019). To ease the comparison across goals with different characteristics, we normalize to one both the horizon and the amount of each goal. For each goal, we then cumulate the amount deposited by the user over time so that we have the percentage of the goal completed as a function of the percentage of the time elapsed. Finally, we compute semi-parametric estimates of the average percentage completion across all goals, as a function of the time elapsed. The black line refers to all goals, the blue line to attained goals and the green line to unattained loans.

Table 1. Summary Statistics

Panel A. Demographic Characteristics

	N	mean	sd	p1	p25	p50	p75	p99
Gender	47,643	0.81	0.39	0.00	1.00	1.00	1.00	1.00
Age	47,216	36.46	11.82	19.00	27.00	35.00	44.00	68.00

Panel B. App Usage by Users

	N	mean	sd	p1	p25	p50	p75	p99
Tenure	27,439	11.48	14.02	1.00	1.00	5.33	16.73	59.27
N. Transactions per Month	27,439	1.14	0.94	0.08	0.57	1.00	1.41	4.23
Net Deposit per Month	27,439	35.04	96.19	-1.84	0.18	5.01	30.00	463.34

Panel C. Accounts and Targets

	N	mean	sd	p1	p25	p50	p75	p99
N. of Savings Accounts	47,746	1.11	0.49	1.00	1.00	1.00	1.00	4.00
Target Present	47,746	0.27	0.44	0.00	0.00	0.00	1.00	1.00
N. Targets	17,240	1.43	1.32	1.00	1.00	1.00	1.00	6.00
N. Active Targets	17,240	0.86	0.67	0.00	1.00	1.00	1.00	4.00
N. Closed Targets	17,240	0.29	0.69	0.00	0.00	0.00	0.00	3.00
N. Won Targets	17,240	0.28	0.97	0.00	0.00	0.00	0.00	4.00

This Table reports cross-sectional summary statistics for the users in our sample. We first compute the value of each variable at the user level and then report the distribution of the variable across all users. For each variable, we report the number of observations used in the second step of the computations, the mean and standard deviation of each variable, as well as the $1st$, $25th$, $50th$, $75th$, and $99th$ percentiles. Panel A refers to the demographic characteristics: *Gender* (1 for males and 0 for females) and *Age* (as of November 2019). Panel B refers to the App usage. The variables in this panel are based only on the individuals that have transacted at least once on the platform. *Tenure* denotes the number of months between the first and the last transaction we observe; *N. Transactions per Month*, denotes the average number of transactions per month (computed as the ratio between the total number of transactions over the full sample and *Tenure*); *Net Deposit per Month* is the average monthly difference between total Deposits and Total Withdrawals made by the user over her *Tenure*. Panel C refers to variables related to the number of accounts and targets associated with each user. *N. of Savings Accounts* is the total number of savings accounts opened by a given user; *Target Present* is the number of accounts per user associated with a target; *N. Targets* is the number of targets; *N. Active Targets*, *N. Closed Targets*, *N. Won Targets* are the number of Active (as of November 2019), Closed and Achieved targets, respectively.

Table 2. Effect of Goal-Setting on Saving Behavior: Baseline Results at the User Level

	(1) Net Deposits	(2) Deposits	(3) Withdrawals
Target Dummy	28.74***	34.56***	5.60***
	(16.37)	(20.24)	(5.67)
User Fixed Effects	✓	✓	✓
Time Fixed Effects	✓	✓	✓
R-Squared	0.34	0.44	0.18
Obs	307,501	307,501	307,501

This table reports coefficient estimates and associated t-statistics (in parentheses) of the following panel regression:

$$Transaction_{i,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,t} + \epsilon_{i,t}$$

where $Transaction_{i,t}$ denotes the monthly amount (in €) transacted by individual i during month t; α_i and α_t identify user and time fixed-effects; $Target\_Dummy$ is an indicator variable equal to "0" if the user has no goals in her account at time t, and equal to "1" if she has at least one goal at time t. The results in the first column are based on $Net\ Deposits$, i.e. the difference between total *Deposits* and $Withdrawals$, while the results in columns (2) and (3) only use *Deposits* and $Withdrawals$, respectively. The standard errors are double-clustered at the user and month levels. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 3. Effect of Goal-Setting on Saving Behavior at the Account Level

Panel A. Specification with User and Time Fixed Effects

	(1) Net Deposits	(2) Deposits	(3) Withdrawals
Target Dummy	26.84***	29.40***	1.61
	(18.31)	(19.57)	(1.61)
User Fixed Effects	✓	✓	✓
Time Fixed Effects	✓	✓	✓
User×Time Effects	✗	✗	✗
R-Squared	0.29	0.39	0.16
Obs	347,411	347,411	347,411

Panel B. Specification with User-Time Fixed Effects

	(1) Net Deposits	(2) Deposits	(3) Withdrawals
Target Dummy	55.16***	56.59***	0.07
	(18.70)	(18.09)	(0.06)
User Fixed Effects	✗	✗	✗
Time Fixed Effects	✗	✗	✗
User×Time Effects	✓	✓	✓
R-Squared	0.56	0.62	0.58
Obs	65,792	65,792	65,792

This table reports coefficient estimates and associated t-statistics (in parentheses) of the following panel regression with user and time fixed effects in Panel A:

$$Transaction_{i,j,t} = \alpha_i + \alpha_t + \beta \; Target\_Dummy_{i,j,t} + \epsilon_{i,j,t}$$

where $Transaction_{i,j,t}$ denotes the monthly amount (in €) transacted by user i in account j at time t; α_i and α_t identify user and time fixed-effects; $Target\_Dummy_{i,j,t}$ is a dummy variable equal to 1 if account j associated with user i at time t is characterized by goal-setting, and zero otherwise. In panel B, we estimate a specification with user-time fixed effects:

$$Transaction_{i,j,t} = \alpha_{i,t} + \beta \; Target\_Dummy_{i,j,t} + \epsilon_{i,j,t}$$

In both panels, the results in the first column are based on *Net Deposits*, i.e. the difference between total *Deposits* and *Withdrawals*, while the results in column (2) and (3) only use *Deposits* and *Withdrawals*, respectively. The standard errors are double-clustered at the user and month levels. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 4. Identification Strategy: Beta-Tester Results

	Panel A. All Users		Panel B. Matched Sample					
			Nearest Neighbor=1		Nearest Neighbor=2		Nearest Neighbor=5	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Treated Dummy	20.01***	16.92***	27.57**	28.68**	19.19**	19.51***	24.57***	22.09***
	(6.75)	(4.22)	(2.92)	(2.97)	(2.74)	(3.52)	(4.93)	(3.41)
User Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Time Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
R-Squared	0.25	0.25	0.35	0.34	0.26	0.25	0.23	0.22
Obs	68,336	74,232	1,999	2,146	2,863	3,068	5,169	5,524

This table reports coefficient estimates and associated t-statistics (in parentheses) of the following differences-in-differences model:

$$Transaction_{i,t} = \alpha_i + \alpha_t + \beta\ Treated_{i,t} + \epsilon_{i,t},$$

where $Transaction_{i,t}$ denotes the monthly net deposit amount (in €) transacted by user i at time t; α_i and α_t identify user and time fixed-effects; $Treated_{i,t}$ is equal to 1 for the beta testers over the beta-testing periods. Because the beta-testing period covers 50 days (from August 1, 2017 to September 20, 2017) and we perform our analysis at the monthly frequency, we adopt two specifications. In the first specification, denoted as (1), the dummy is equal to 1 only in August 2017, while in the second specification, denoted as (2), the dummy is equal to 1 in both August and September 2017. The estimation window ranges from August 2016 to August 2017 in (1) and from August 2016 to September 2017 in (2). Panel A uses all the users excluded from being beta testers as controls. In Panel B, the control group comprises non-beta testers matched to the characteristics of the beta-testers according to the following characteristics computed from August 2016 to July 2017: *Male*, whether the user is male; *Age*; *Account Num*, the number of accounts associated with each user; *Net Inflow*, the monthly net deposits into the account; *tenure*, the number of months the user has used the app as July 2017; and *Risk*, the average risk of the funds selected by the user. The results focus on nearest neighbor matches that associate each beta-tester to a single user on the App in the first two columns, to the closest two matches in the middle columns and the closest five matches in the last two columns. The standard errors are double-clustered at the user and month levels. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 5. Statistics of Matched Sample

Panel A: Beta Testers

	Mean	Std	p5	p25	p50	p75	p95
Tenure	516.91	309.35	149.00	304.00	404.00	668.00	1,156.00
N. Accounts	1.00	0.00	1.00	1.00	1.00	1.00	1.00
Net Inflow	181.44	245.01	6.03	22.43	68.74	254.59	676.41
Risk	8.66	3.44	3.25	6.19	7.55	11.12	16.36
Gender	0.86	0.35	0.00	1.00	1.00	1.00	1.00
Age	41.21	9.62	26.00	34.00	41.00	48.00	58.00

Panel B: Other Users

	t-test	Mean	Std	p5	p25	p50	p75	p95
Tenure	-2.68	430.76	319.38	32.00	183.00	345.00	628.00	1,117.00
N. Accounts	0.20	1.00	0.06	1.00	1.00	1.00	1.00	1.00
Net Inflow	-1.85	134.35	265.37	1.01	7.88	30.00	105.62	733.33
Risk	-1.70	8.05	3.49	2.22	6.19	7.55	9.45	15.40
Gender	-1.49	0.81	0.39	0.00	1.00	1.00	1.00	1.00
Age	-0.31	40.90	11.60	24.00	32.00	40.00	49.00	61.00

Panel C: Matched Users

	t-test	Mean	Std	p5	p25	p50	p75	p95
Tenure	0.25	528.78	331.10	45.00	263.00	455.00	814.00	1,184.00
N. Accounts	0.00	1.00	0.00	1.00	1.00	1.00	1.00	1.00
Net Inflow	-0.12	176.48	289.45	1.01	11.00	50.28	171.36	1,000.00
Risk	-0.69	8.29	3.87	2.22	6.19	7.55	9.45	16.36
Gender	0.43	0.88	0.32	0.00	1.00	1.00	1.00	1.00
Age	-0.21	40.87	11.94	24.00	33.00	39.00	48.00	65.00

This Table presents cross-sectional summary statistics for the sample of beta testers (Panel A), the rest of the users (Panel B) and the matched users (Panel C). We first compute the value of each variable at the user level and then report the distribution of the variable across all users. We report the mean and standard deviation of each variable, as well as the $5th$, $25th$, $50th$, and $75th$, and $95th$ percentile. $Tenure$ is the number of months elapsed from the opening of user's first account to July 2017—right before the beginning of the beta-testing period; *N. Accounts*, is the number of accounts opened from August 2016 to July 2017; $Net\ Inflow$ is the monthly average net deposit into the App from August to July 2017; *Risk*, is the average annual standard deviation across the mutual funds used by the user; *Gender*, is a dummy variable equal to one if male; and *Age* is age measured as of July 2017. In Panels B and C we also report the t-stat associated with tests on the equality of means between beta-testers and non-beta-testers along the various characteristics. To create the sample of users in Panel C, we use a nearest neighbor propensity score matching procedure whereby each beta-tester is matched to a single user on the App.

Table 6. Heterogeneity in Target Characteristics

	Target Type	Target Horizon	Target Size
Target Dummy×Interaction_1	-4.57**	6.02***	-2.01
	(-2.63)	(3.21)	(-1.13)
Target Dummy×Interaction_2	-3.05		
	(-1.06)		
Target Dummy×Interaction_3	-8.52***		
	(-3.53)		
Target Dummy	29.68***	23.56***	27.90***
	(17.90)	(12.18)	(13.18)
User Fixed Effects	✓	✓	✓
Time Fixed Effects	✓	✓	✓
R-Squared	0.29	0.29	0.29
Obs	347,331	347,411	347,411

This table reports coefficient estimates and associated t-statistics (in parentheses) of the following panel regression:

$$M\_Net\_Deposits_{i,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,t} + \sum_{k=1}^{3} \gamma_k\ Target\_Dummy_{i,t} \times Interaction\_k_i + \epsilon_{i,t}$$

where $M\_Net\_Deposits_{i,t}$ denotes the monthly net amount (in €) deposited by user i at time t; α_i and α_t identify user and time fixed-effects; $Target\_Dummy_{i,t}$ is a dummy variable equal to 1 if user i at time t has a target and $Interaction\_1_i$ is a dummy equal to 1 if the target falls in the leisure category and zero otherwise, $Interaction\_2_i$ is a dummy equal to 1 if the target falls in the durable category and zero otherwise and $Interaction\_3_i$ is a dummy equal to 1 if the target falls in the "other" category and zero otherwise. In the second column, $Interaction\_1_i$ identifies targets with horizons of less than (or equal to) a year and in the third column $Interaction\_1_i$ identifies targets smaller than €1,000. The standard errors are double-clustered at the user and month levels. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 7. Determinants of Goal Achievement

	Spec 1	Spec 2	Spec 3	Spec 4
Hobby	-0.54***	-0.56***	-0.11***	-0.11***
	(-7.12)	(-7.13)	(-7.75)	(-7.73)
Travel	-0.75***	-0.73***	-0.13***	-0.13***
	(-13.26)	(-12.68)	(-14.98)	(-14.30)
Other	-0.53***	-0.50***	-0.10***	-0.09***
	(-8.00)	(-7.26)	(-8.73)	(-7.98)
Car	-0.83***	-0.82***	-0.12***	-0.11***
	(-8.60)	(-8.27)	(-9.70)	(-9.07)
House	-0.62***	-0.63***	-0.11***	-0.11***
	(-6.56)	(-6.52)	(-7.27)	(-7.35)
Target Size	-0.68***	-0.72***	-0.13***	-0.13***
	(-28.01)	(-28.31)	(-28.60)	(-28.80)
Horizon	-0.11***	-0.17***	-0.02***	-0.03***
	(-4.72)	(-6.90)	(-5.03)	(-7.05)
Male	-0.45***	-0.47***	-0.07***	-0.08***
	(-7.62)	(-7.81)	(-8.15)	(-8.55)
Age	0.47***	0.45***	0.09***	0.08***
	(22.74)	(20.73)	(23.42)	(21.08)
Constant	-0.61***	-4.72***	0.37***	0.11
	(-20.37)	(-3.68)	(61.24)	(0.81)
Time FE	✗	✓	✗	✓
Province FE	✗	✓	✗	✓
R-Squared	0.13	0.16	0.14	0.18
Obs	13,838	13,823	13,838	13,838

This table reports coefficient estimates and associated t-statistics (in parentheses) of the following regression:

$$p_i = Pr\left(y_i = 1|\boldsymbol{x}_i\right) = \frac{\exp\left(\boldsymbol{x}_i'\boldsymbol{\beta}\right)}{1+\exp\left(\boldsymbol{x}_i'\boldsymbol{\beta}\right)},$$

where the dependent variable is 1 if goal i is achieved and 0 otherwise. The vector $\mathbf{x}_i$ includes dummies for the following categories of goals: hobby, travel, car, house, and generic savings (the base case), and controls including the log transformation of the goal horizon, measured in days; the log transformation of the goal amount; the age of the user; a male dummy; monthly and province fixed-effects. The first and second (third and fourth) columns report coefficients for logit (linear probability) models. The second and fourth column include monthly time effects and province fixed-effects. All the continuous variables have been standardized so that they have a standard deviation of one. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table 8 Goal-Setting and Other Undersaving Reasons

		Financial Literacy	Patience	Education	Income	Attention
Target Dummy	33.98***	31.94***	34.60***	33.54***	34.70***	32.05***
	(7.45)	(6.24)	(7.55)	(6.82)	(6.34)	(11.16)
Target Dummy×Inter		19.82**	2.38	6.67	1.16	-9.53*
		(2.27)	(0.22)	(0.74)	(0.13)	(1.84)
User Fixed Effects	✓	✓	✓	✓	✓	✓
Time Fixed Effects	✓	✓	✓	✓	✓	✓
R-Squared	0.34	0.33	0.33	0.33	0.33	0.32
Obs	19,279	17,698	17,698	17,577	17,698	226,222

This table reports coefficient estimates and associated t-statistics (in parentheses) of the following panel regression:

$$M\_Net\_Deposits_{i,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,t} + \gamma\ Target\_Dummy_{i,t} \times Inter_i + \epsilon_{i,t}$$

where $M\_Net\_Deposits_{i,t}$ denotes the monthly net amount (in €) deposited by user i at time t; α_i and α_t identify user and time fixed-effects; $Target\_Dummy_{i,t}$ is a dummy variable equal to 1 if user i at time t has a target. The first column does not include any interaction and simply re-estimates the specification in column 1 of Table 2 on the sample of survey respondents. Columns two through six focus on financial literacy, patience education, income and users' number of logins before implementing goal setting, respectively. Refer to Section 5.2 for their definition. The standard errors are double-clustered at the user and month levels. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Online Appendix to
Goal Setting and Saving in the FinTech Era

Antonio Gargano
University of Houston

Alberto G Rossi
Georgetown University

(Not for publication)

Appendix A Likelihood of Achieving Goals: Alternative Categorization of Goals

We use data from the text entered by the users to describe the goals and manually re-categorize the goals using both a fine and a coarse scheme. The fine categorization includes *Travel*, *Vehicle*, *House*, *Celebrations*, *Family*, *Hobby*, *Consumer Goods*, *Health*, *Liability*, *Entrepreneurship*, *Saving* and *Other*.[16] The coarse categorization matches the one available on the App.[17] We also use a machine learning technique, the Latent Dirichlet Allocation (LDA), to create topics and assign targets to them. Consistent with the manual classification, we impose a fine scheme (which includes 11 topics in addition to the general *Savings* category) and a coarse scheme (which includes 5 topics in addition to the general *Savings* category). The advantage of this approach is that the resulting categorization does not involve human judgement.

We then re-estimate the linear probability model in Equation (6) using dummies based on both categorizations and keeping *Savings* as our base case. If selecting a more specific goal improves the achievement rate but the pre-specified categories available on the App do not fully span the purpose of the goals set by the users, we would expect at least one of the finer categories to be associated with a positive coefficient. Similarly, if the users did not pay attention when selecting the goal category, we should find results different from Table 7 using the coarser manual categorization. We present the results in Table A.4, where columns 1 and 2 report the results for the fine and coarse manual classification, while columns 3 and 4 report the results for the machine learning classification.

Starting from the manual categorization, we find again that specific goals are negatively related to the probability of success. Relative to the category used as base (general *Savings*), the targets in the fine categories have a lower probability of being achieved. The *Health* category displays the largest

[16] The *Holiday* goal includes words like travel, holiday, trip, cruise, flight or geographical locations; the *Vehicle* goal includes words like car, motorbike, scooter or their models/manufacturers; the *House* goal includes words like house, apartment, renovation, and specific furniture and appliances; the *Celebrations* goal includes words like wedding, baptism and communion; the *Family* goals includes words like daughter, son, kids, baby, and the Italian first names; the *Consumer Goods* goals includes words like iPhone, pc, camera or brands/manufacturers; the *Hobby* goals includes words like tennis, surfing and other leisure activity; the *Health* goal includes words like health, dentist, and surgery; the *Liability* goal includes words like debt, taxes and repayment; the *Entrepreneurship* goal includes words like start, project, and activity; the general *Saving* goal includes words like saving, income and investment without any further characterization.

[17] We do so by merging *Celebrations*, *Family*, *Consumer Goods*, *Health*, *Liability*, and *Entrepreneurship* into the *Other* category.

difference (-18 percentage points), while *Family* displays the smallest one (-5 percentage points). When we use the coarse scheme, which matches the pre-specified categories available on the App, we obtain point estimates almost identical to the last column of Table 7. The classification based on the LDA algorithm gives virtually identical results.

These results underscore the robustness of our finding that precise goals are achieved less often than generic goals.

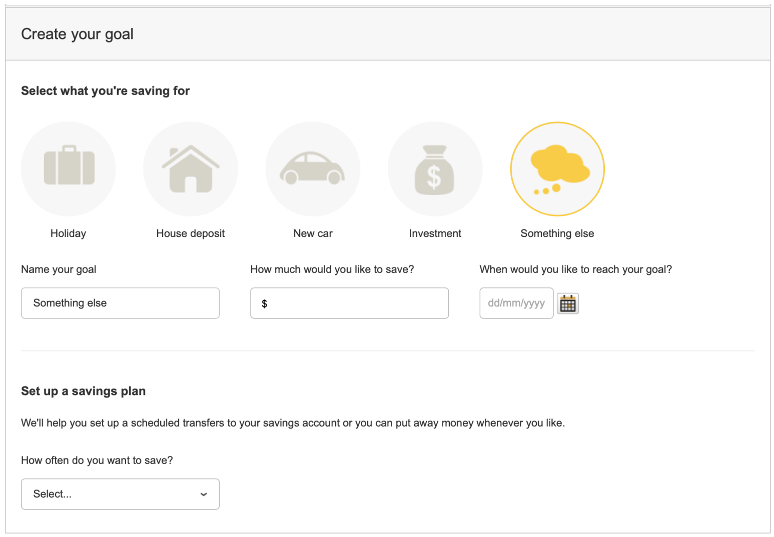

(a) Commonwealth Bank of Australia

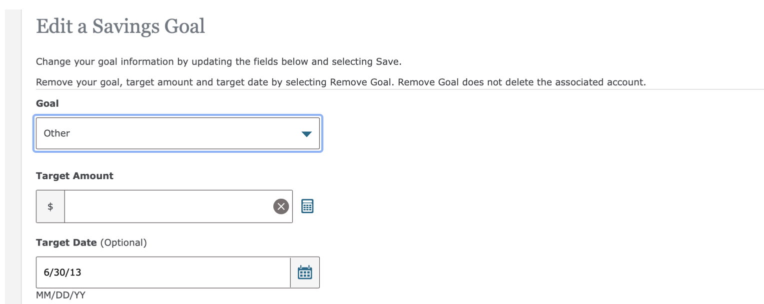

(b) Wells Fargo

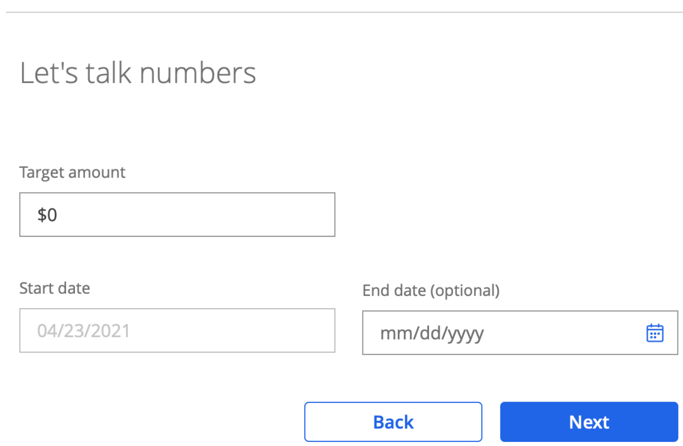

(c) Chase Bank

Figure A.1: This Figure shows examples of goal setting for saving App implemented by Commonwealth Bank of Australia (a), Wells Fargo (b), and Chase Bank (c).

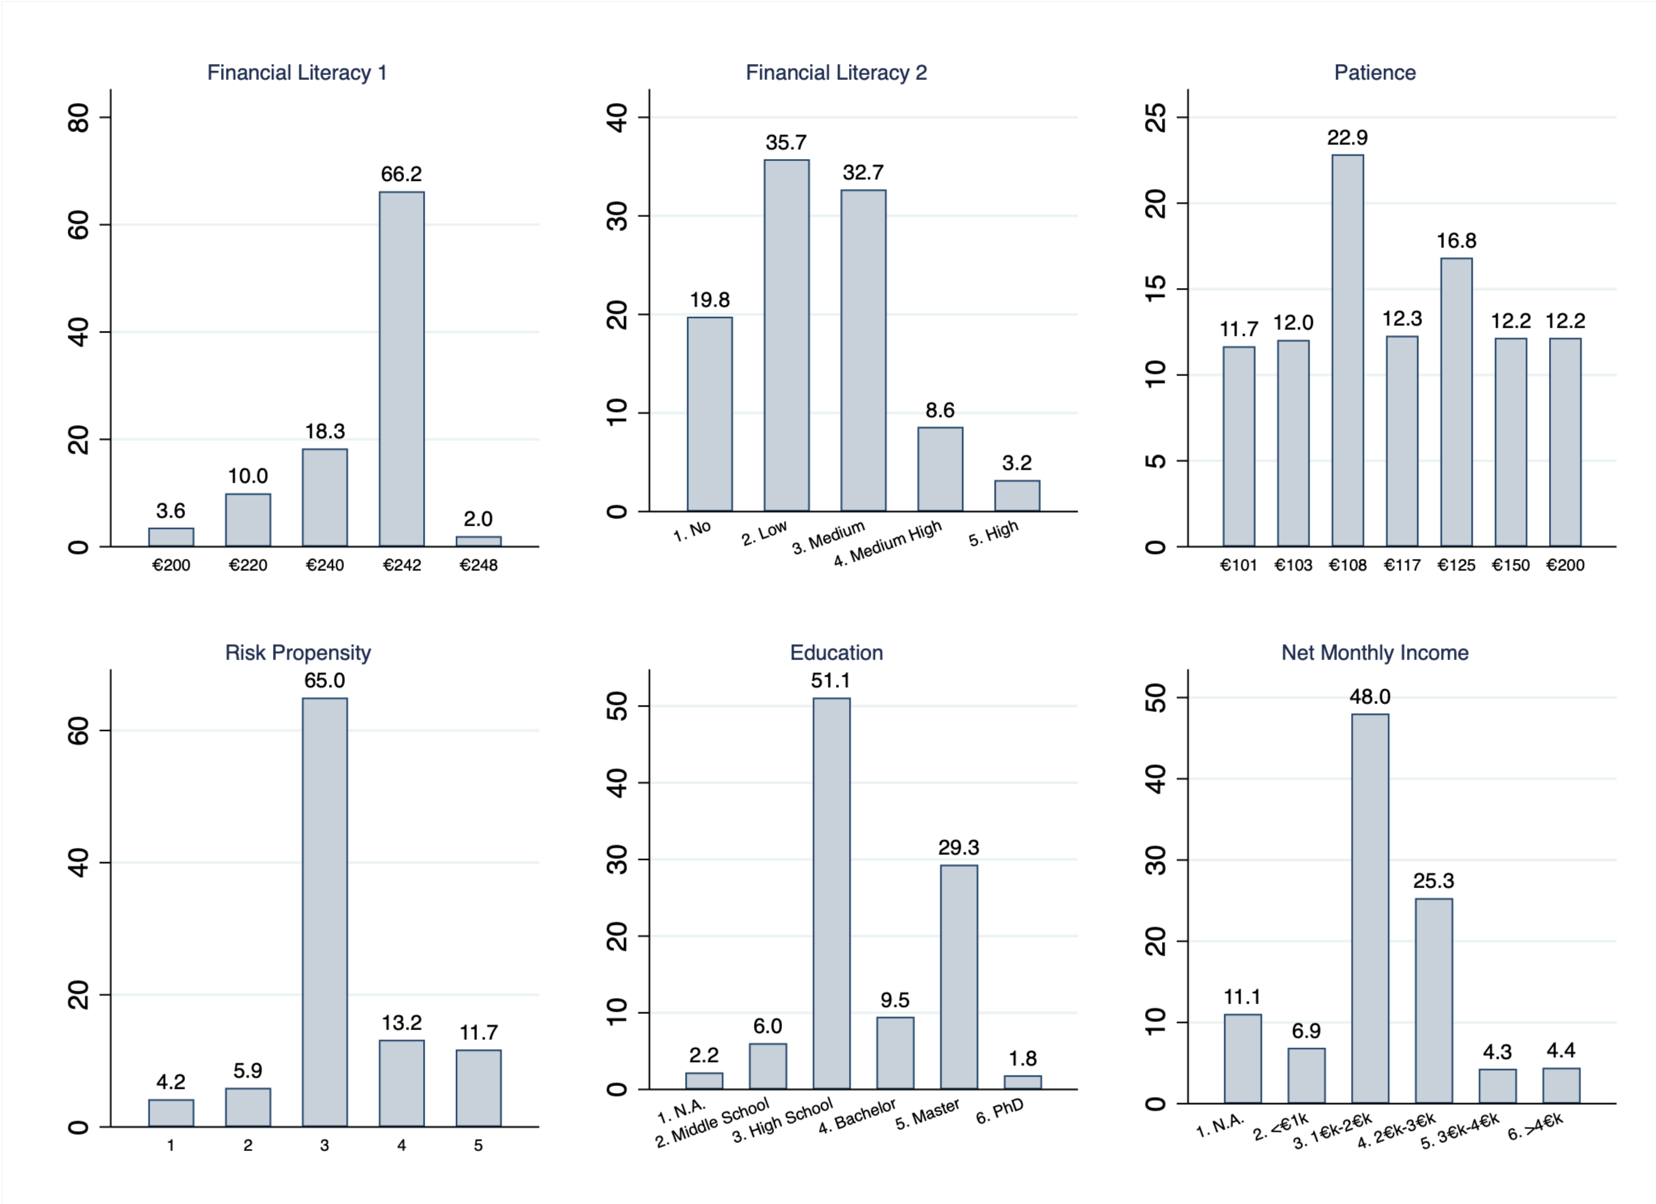

Figure A.2: This Figure shows the distribution of the answers to the survey questions relative to users' characteristics and economic preferences. The question in the upper-left plot is *"Let's say you have $200 in a savings account. The account earns 10% interest per year. If you never withdraw money or interest payments, how much will you have in the account at the end of 2 years?"* The question in the upper-middle plot is *"How knowledgeable are you about investments and financial markets?"* The question in the upper-right plot is *"Imagine, you get either $100 immediately or a higher amount of money in a month. What is the lowest amount you would be willing to wait for a month?"* The question in the lower-left plot is *"On a scale from 1 to 5, how would you rate your willingness to take risks regarding financial matters?"* The question in the lower-middle plot is *"What is your level of education?"* The question in the lower-right plot is *"What is your monthly net income?"* We display the percentage of respondents on top of each bar.

Table A.1. Fund Performance

Panel A: Fixed Income Funds

	Mean	Median	Volatility	Sharpe
12-months	0.30	-0.00	0.33	0.93
Short Term	2.43	1.81	1.84	1.32
Euro Bond	2.86	3.02	2.75	1.04
Performance	2.10	0.42	8.66	0.24

Panel B: Equity Funds

	Mean	Median	Volatility	Sharpe
America	5.71	5.04	13.90	0.41
Asia-Pacific	2.83	10.24	11.80	0.24
Europe	4.21	9.63	14.28	0.29
Active ETF	1.67	3.73	10.77	0.15
Global	4.13	7.25	9.75	0.42
Italy	0.17	15.64	17.79	0.01
Emerging Markets	4.75	8.77	12.67	0.37

Panel C: Balanced Funds

	Mean	Median	Volatility	Sharpe
Aggressive	2.54	4.08	6.90	0.37
Dynamic	2.28	3.39	5.56	0.41
Immune	0.74	0.00	1.73	0.43
Cautious	1.59	-0.37	8.40	0.19

This table presents summary statistics of the returns characteristics of the mutual funds available on the App. The funds are categorized in Fixed Income funds (Panel A), Equity Funds (Panel B) and Balanced Funds (Panel C). For each fund, we report the average and median of the annualized returns, annualized realized volatility—computed using daily returns—and annualized Sharpe ratios.

Table A.2. Statistics of Matched Samples (Robustness)

Panel A: Beta Testers

	Mean	Std	p5	p25	p50	p75	p95
Tenure	516.91	309.35	149.00	304.00	404.00	668.00	1,156.00
N. Accounts	1.00	0.00	1.00	1.00	1.00	1.00	1.00
Net Inflow	181.44	245.01	6.03	22.43	68.74	254.59	676.41
Risk	8.66	3.44	3.25	6.19	7.55	11.12	16.36
Gender	0.86	0.35	0.00	1.00	1.00	1.00	1.00
Age	41.21	9.62	26.00	34.00	41.00	48.00	58.00

Panel B: Matched Sample (2 Nearest Neighbors)

	t-test	Mean	Std	p5	p25	p50	p75	p95
Tenure	0.23	526.02	326.50	66.00	275.00	457.00	795.00	1,184.00
N. Accounts	0.00	1.00	0.00	1.00	1.00	1.00	1.00	1.00
Net Inflow	-0.38	168.42	305.44	1.01	10.00	50.00	155.69	1,000.00
Risk	-0.49	8.43	3.99	2.22	6.19	7.55	9.45	16.36
Gender	0.46	0.88	0.32	0.00	1.00	1.00	1.00	1.00
Age	-0.45	40.62	11.95	24.00	30.00	40.00	49.00	61.00

Panel C: Matched Sample (5 Nearest Neighbors)

	t-test	Mean	Std	p5	p25	p50	p75	p95
Tenure	-0.08	513.97	325.39	85.00	262.00	438.00	775.00	1,184.00
N. Accounts	0.00	1.00	0.00	1.00	1.00	1.00	1.00	1.00
Net Inflow	-0.61	163.73	298.65	1.01	10.00	49.31	135.85	1,000.00
Risk	0.24	8.76	3.98	2.22	6.19	7.55	9.66	16.36
Gender	-0.06	0.86	0.35	0.00	1.00	1.00	1.00	1.00
Age	-1.02	40.05	11.59	24.00	31.00	39.00	47.00	60.00

This Table presents cross-sectional summary statistics for the sample of beta testers (Panel A), and two samples of matched users (Panels B and C). We first compute the value of each variable at the user level and then report the distribution of the variable across all users. We report the mean and standard deviation of each variable, as well as the $5th$, $25th$, $50th$, and $75th$, and $95th$ percentile. $Tenure$ is the number of months elapsed from the opening of user's first account to July 2017—right before the beginning of the beta-testing period; *N. Accounts*, is the number of accounts opened from August 2016 to July 2017; $Net\ Inflow$ is the monthly average net deposit into the App from August 2016 to July 2017; *Risk*, is the average annual standard deviation across the mutual funds used by the user; *Gender*, is a dummy variable equal to one if male; and *Age* is age measured as of July 2017. In Panels B and C we also report the t-statistics associated with tests on the equality of means between beta-testers and non-beta-testers along the various characteristics. To create the sample of users in Panel B (C), we use a nearest neighbor propensity score matching procedure whereby each beta-tester is matched to the closest 2 (5) matches on the App.

Table A.3. Effect of Goal-Setting on Saving Behavior:
Assessing the Importance of Saving Behavior Outside the App.

	Baseline	Baseline No Savings Outside	No Savings Outside	No Change in Saving Outside	Stop Saving Outside
Target Dummy× Interaction			4.61	-3.09	13.01
			(0.41)	(-0.37)	(1.24)
Target Dummy	33.98***	37.26***	34.00***	35.40***	32.27***
	(7.45)	(3.80)	(7.14)	(6.00)	(6.50)
User Fixed Effects	✓	✓	✓	✓	✓
Time Fixed Effects	✓	✓	✓	✓	✓
R-Square	0.34	0.46	0.33	0.34	0.34
Obs	19,279	2,002	19,027	19,279	19,279

This table reports coefficient estimates and associated t-statistics (in parentheses) of the following panel regressions in columns 1 and 2:

$$Transaction_{i,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,t} + \epsilon_{i,t}$$

where $Transaction_{i,t}$ denotes the monthly amount (in €) transacted by individual i during month t; α_i and α_t identify user and time fixed-effects; $Target\_Dummy$ is an indicator variable equal to "0" if the user has no goals in her account at time t, and equal to "1" if she has at least one goal at time t. The results in column 1 are computed using all survey respondent data. The results in column 2 use only the survey respondents that state in the survey they have no savings outside the App. In Columns 3 through 5 we estimate:

$$Transaction_{i,t} = \alpha_i + \alpha_t + \beta\ Target\_Dummy_{i,t} + \gamma\ Target\_Dummy_{i,t} \times Interaction_i + \epsilon_{i,t}$$

where all regressors are defined as in the previous regression specification and the only difference stands in the variable $Interaction_i$. In column 3, $Interaction_i$ is equal to 1 for those individuals who respond they have no savings outside the App and is equal to zero otherwise. In column 4, $Interaction_i$ is equal to 1 for those individuals who respond they do not change their saving rate outside of the App after adopting the App and is equal to zero otherwise. In column 5, $Interaction_i$ is equal to 1 for those individuals who respond they reduce their savings outside the App after adopting the App and is equal to zero otherwise. The standard errors are double-clustered at the user and month levels. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Table A.4. Determinants of Goal Achievement: Robustness

	Panel A: Manual			Panel B: Machine Learning	
	Fine	Coarse		Fine	Coarse
Hobby	-0.07***	-0.07***	Topic 1	-0.11***	-0.11***
	(-2.81)	(-2.81)		(-11.88)	(-11.45)
Travel	-0.14***	-0.13***	Topic 2	-0.11***	-0.12***
	(-13.36)	(-13.30)		(-8.03)	(-9.53)
Other	-0.06***	-0.08***	Topic 3	-0.11***	-0.10***
	(-4.30)	(-7.84)		(-7.77)	(-6.46)
Car	-0.13***	-0.13***	Topic 4	-0.11***	-0.10***
	(-9.58)	(-9.51)		(-5.80)	(-4.89)
House	-0.11***	-0.10***	Topic 5	-0.07***	-0.09***
	(-6.20)	(-6.17)		(-3.14)	(-3.54)
Celebration	-0.08***		Topic 6	-0.09***	
	(-4.30)			(-2.76)	
Consumer Goods	-0.12***		Topic 7	-0.08**	
	(-7.70)			(-2.09)	
Entrepreneurship	-0.11**		Topic 8	-0.09*	
	(-2.43)			(-1.72)	
Family	-0.05**		Topic 9	0.06	
	(-2.12)			(1.01)	
Liability	-0.13***		Topic 10	-0.06	
	(-3.45)			(-0.58)	
Health & Wellbeing	-0.18***		Topic 11	-0.20***	
	(-4.95)			(-3.54)	
Constant	0.02	0.01		-0.00	-0.02
	(0.17)	(0.08)		(-0.01)	(-0.13)
Controls	✓	✓		✓	✓
Time FE	✓	✓		✓	✓
Province FE	✓	✓		✓	✓
R-Square	0.18	0.18		0.18	0.18
Obs	12,940	12,940		12,940	12,940

This table reports coefficient estimates and associated t-statistics (in parentheses) of the following linear probability model:

$$p_i = Pr\left(y_i = 1 | \boldsymbol{x}_i\right) = \boldsymbol{x}_i'\boldsymbol{\beta} + \boldsymbol{c}_i'\boldsymbol{\gamma} + \alpha_{time} + \alpha_{province} + \epsilon_i$$

where the dependent variable is 1 if goal i is achieved and 0 otherwise. The vector $\mathbf{x}_i$ contains the dummies for the categories of goals described in Section Appendix A. In all specifications, the generic *Saving* category is the base case. The set of controls ($\boldsymbol{c}$) include the log transformation of the goal horizon, measured in days; the log transformation of the goal amount; the age of the user; and a gender dummy. α_{time} and $\alpha_{province}$ denote monthly and province fixed-effects, respectively. All the continuous variables have been standardized to have a standard deviation of one. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level.

Section 1 of 4: Saving Behavior Outside Gimme5

Q1.1 **Since using** Gimme5, how much have you saved on average per month **outside the App**?

- €0
- Between €1 and €49
- Between €50 and €100
- Between €101 and €150
- More than €150
- Do not know / Do not remember

Q1.2 **Before using** Gimme5, how much have you saved on average per month **outside the App**?

- €0
- Between €1 and €49
- Between €50 and €100
- Between €101 and €150
- More than €150
- Do not know / Do not remember

Q1.3 **Since using** Gimme5, have you changed how much you saved **outside the App?**

- Yes, decreased
- No
- Yes, increased
- Do not know / Do not remember

Q1.4 Which channels do you use to invest your savings **outside the App?**
(please select one or more)

- I do not save outside the App
- Bank Saving Account
- Trading Platform
- Financial Advisor
- Other Apps
- Other ____

Section 2 of 4: Usage of Saving Goals

Q2.0.1 How many saving goas have you created so far?

- 0
- 1
- 2
- 3
- 4
- 5 or more

-- Note: Questions Q.2.1.1 to Q.2.1.3 below are only visible for those who chose 0 in Q2.0.1 --

Q2.1.1 Why have not created a saving goal yet?

- I did not know
- I'm not interested in saving for a specific goal
- It is too difficult
- I do not know
- Other ____

Q2.1.2 If you decided to create a saving goal, what aspect would you find more motivating?

- Having a deadline
- The purpose of the goal
- The ability to monitor my progress towards the goal
- I do not know
- Other ____

Q2.1.3 What decision associated with creating a saving goal do you find more challenging?

- Deciding the amount
- Deciding the purpose of the goal
- Deciding the horizon
- I do not know
- Other ____

--- Note: Questions Q.2.2.1 to Q.2.2.4 below are only visible for those who chose 1, 2, 3, 4, 5 or more in Q2.0.1 --

Q2.2.1 In your experience, what is the most motivating aspect of having a saving goal?

- Having a deadline
- The purpose of the goal
- The ability to monitor my progress towards the goal
- Other ____
- I do not think having a goal is motivating

Q2.2.2 Think about the saving goals you **have achieved.** In your opinion, what has been the most decisive factor?

- The deadline
- The purpose of the goal
- The ability to monitor my progress towards the goal
- The return of the fund
- I have never achieved my saving goals
- I do not know
- Other ____

Q2.2.3 Think about the saving goals you **have NOT achieved.** In your opinion, what has been the most decisive factor?

- Lack of initial motivation
- The amount of the goal was too high compared to the horizon
- I faced unexpected expenses
- The return of the fund
- I have always achieved my saving goals
- I do not know
- Other ____

Q2.2.4 Since you started using saving goals your use of Joinks

- Is decreased a lot
- Is decreased a little
- Is remained constant
- Is increased a little
- Is increased a lot
- Do not know
- I do not use Joinks

--- Note: Questions Q.2.3.1 to Q.2.3.3 below are visible to all respondents --

Q2.3.1 What feature of the App do you find most helpful in increasing your saving?

- Saving small amounts through Joinks
- Investing in mutual funds
- Saving goals
- Automatic saving rules
- Do not know
- Other ____

Q2.3.2 How can we help you to achieve your goals?
(Please select one or more)

- Suggesting the horizon of the goal
- Suggesting the amount of the goal
- Suggesting how much to save per month based on the horizon and the amount of the goal
- Sending a monthly reminder on how much to save
- Do not know
- Other ____

Q2.3.3 What innovations would you like to be implemented in the next release of the App?
(Specify below or leave empty)

____

____

____

Section 3 of 4: Usage of the Joink Function

Q3.0.1 Do you find the Joink Function easy to understand?

- Yes
- No
- Do not know

Q3.0.2 On average, how many times in a month do you use the Joink function?

- Never
- Between 1 and 5
- Between 6 and 10
- More than 10

--- Note: Question Q3.1.1 below is only visible for those who chose "Never" in Q3.01. --

Q3.1.1 Why do you not use the Joink function?

- I do not know what it is
- I do not know how to use it
- I do not think it is useful
- I do not now
- Other ____

--- Note: Questions Q3.2.1 and Q3.2.2 below are only visible for those who did not choose "Never" in Q3.01. --

Q3.2.1 When do you use the Joink Function?

- When I give up a daily expense
- When I perform a daily expense
- When I give up an occasional expense
- When I perform an occasional expense
- My use of the Joink function is unrelated to any specific situation
- Other ____

Q3.2.2 What is the main reason you use Joinks?
(Specify below or leave empty)

____

____

____

Section 4 of 4: Your Opinion on Financial Markets

Q4.1 Let's say you have $200 in a savings account. The account earns 10% interest per year. If you never withdraw money or interest payments, how much will you have in the account at the end of 2 years?

- €200
- €220
- €240
- €242
- €248

Q4.2 How knowledgeable are you about investments and financial markets?

- Not an expert
- Low
- Medium
- Medium High
- High

Q4.3 Imagine you get either $100 immediately or a higher amount of money in a month. What is the lowest amount you would be willing to wait for a month?

- €101
- €103
- €108
- €117
- €125
- €150
- €200

Q4.4 On a scale from 1 to 5, how would you rate your willingness to take risks regarding financial matters?

- 1, I avoid risk at all costs
- 2
- 3, I try to balance risk and returns
- 4
- 5, I want the highest possible returns even if it entails a high level of risk

Q4.5 What is your highest level of education?

- Middle School
- High School
- Bachelor
- Master
- PhD
- Prefer not to answer

Q4.6 What is your net monthly income?

- Less than €1.000
- Between €1.000 and €2.000
- Between €2.000 and €3.000
- Between €3.000 and €4.000
- More than €4.000
- Prefer not to answer