

Dur, Robert; Fleming, Dimitry; van Garderen, Marten; van Lent, Max

Working Paper

A Social Norm Nudge to Save More: A Field Experiment at a Retail Bank

CESifo Working Paper, No. 8894

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Dur, Robert; Fleming, Dimitry; van Garderen, Marten; van Lent, Max (2021) : A Social Norm Nudge to Save More: A Field Experiment at a Retail Bank, CESifo Working Paper, No. 8894, Center for Economic Studies and Ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/232491>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

A Social Norm Nudge to Save More: A Field Experiment at a Retail Bank

Robert Dur, Dimitry Fleming, Marten van Garderen, Max van Lent

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

A Social Norm Nudge to Save More: A Field Experiment at a Retail Bank

Abstract

A large fraction of households have very little savings buffer and are therefore vulnerable to financial shocks. This paper examines whether a social norm nudge can induce such households to save more. We ran a large-scale field experiment at a retail bank in the Netherlands. We find that households who are exposed to the social norm nudge click more often on a link to a personal web page where they can start or adjust an automatic savings plan. However, analyzing detailed bank data, we find no treatment effect on actual savings, neither in the short run nor in the long run. Our null findings are quite precisely estimated. A complementary small-scale survey experiment suggests that people did notice the social norm nudge and also that it had an impact on savings intentions.

JEL Codes: C930, D140, D900, E210, G400.

Keywords: household savings, field experiment, nudges, social norms.

Robert Dur
Department of Economics
Erasmus University Rotterdam / The Netherlands
dur@ese.eur.nl

Marten van Garderen
ING Economics Bureau
ING Netherlands / Amsterdam
Marten.van.Garderen@ing.com

Dimitry Fleming
ING Economics Bureau
ING Netherlands / Amsterdam
dimitry.fleming@ing.com

Max van Lent
Department of Economics
Leiden University / The Netherlands
m.van.lent@law.leidenuniv.nl

February 6, 2021

1 Introduction¹

A large fraction of households have very little savings buffer and are therefore vulnerable to financial shocks. For instance, in the US more than 1 out of 4 households have hardly any liquid savings (Bhutta and Dettling 2018). The same holds for 40 percent of the working-age population in the UK.² In the Netherlands, 1 out of 3 households have a buffer that is too low according to the Dutch Institute for Budgetary Research and Education (Nibud 2017). Households with too little savings buffer are at risk of having to take up expensive loans and defaulting.

Inducing households to increase savings has turned out to be a major challenge. Interventions that provide financial education or information have often failed to create substantial and lasting behavior change or are very expensive.³ A low cost intervention that has proven to be successful in many other settings is social norm nudging: informing people that their behavior deviates from what most others do

¹This report is based on anonymized data from customers of ING Netherlands. Data was treated in strict compliance with the General Data Protection Regulation. The report has been prepared by the authors for the TFI long-term research track. The views and opinions expressed in this report are solely those of the authors and do not necessarily reflect the official policy or position of the Think Forward Initiative – TFI – or any of its partners. Responsibility for the data analyses and content in this report lies entirely with the authors. The primary purpose of the TFI Research Programme is to inspire practical research insights in the financial decision-making domain. It does not constitute any financial advice or service offer.

The data used in this study are confidential and cannot be shared publicly.

²See the data from Money Advice Service released in 2016: <https://www.moneyadviceservice.org.uk/en/corporate/press-release-low-savings-levels-put-millions-at-financial-risk>.

³See e.g. Bernheim and Garrett (2003), Lusardi (2004), Bell et al. (2008), Skimmyhorn (2016), Brown et al. (2016), and Urban et al. (2020). Beshears et al. (2018) provide a recent review of the literature.

has been found to be a powerful trigger to change behavior in the direction of the descriptive social norm.⁴ Social norm nudging has so far been rarely studied as a way to encourage households to increase their savings.⁵ Social norm nudges hold promise in this context given the abundance of evidence on peer effects in financial decisions.⁶

We set up a large-scale field experiment at a retail bank in the Netherlands (ING Netherlands) to study the effect of a social norm nudge on households' savings behavior. Our social norm nudge targets households whose savings buffer is less than that of the median household in their neighborhood. We examine the effect of the message: "You have a lower buffer with us than most other ING clients in your neighborhood".⁷ We include the nudge in an email, sent by ING Netherlands in January 2018, that intends to promote savings by households. We estimate the causal effect of the nudge by comparing savings of households that received the email

⁴See e.g. Wechsler et al. (2003), Frey and Meier (2004), Schultz et al. (2007), Goldstein et al. (2008), Gerbers and Rogers (2009), Chen et al. (2010), Allcott (2011), Allcott and Rogers (2014), Bradler et al. (2016), Hallsworth et al. (2016), Coffman et al. (2017), Hallsworth et al. (2017), Brandon et al. (2017), Bhanot (2018), Giaccherini et al. (2019), and Bott et al. (2020). However, norm-nudge interventions are not always successful, see among others Blumenthal et al. (2001), Fellner et al. (2013), John and Blume (2018), Cranor et al. (2020), and Dimant et al. (2020).

⁵The only studies we are aware of are Beshears et al. (2015) and Kast et al. (2018). We discuss how we relate to these studies at the end of this section.

⁶See e.g. Duflo and Saez (2002, 2003), Hong et al. (2004), Brown et al. (2008), Kuhn et al. (2011), Brown and Laschever (2012), Bursztyn et al. (2014), Lieber and Skimmyhorn (2018), and Ouimet and Tate (2020).

⁷Social norm nudges that use households in the neighborhood as peer group have been effective in reducing energy consumption (Schultz et al. 2007, Allcott 2011, Allcott and Rogers 2014, and Brandon et al. 2017). Giaccherini et al. (2019) study the effect of a social norm nudge on household purchase of energy efficient lightbulbs, varying the proximity of the reference group in the social norm nudge. They find that a social norm nudge that compares the consumer with other households "in this area" performs much better than a social norm nudge that compares the consumer with other U.S. households. Note also that some of the studies mentioned in footnote 6 find evidence for neighborhood peer effects in financial decisions (Hong et al. 2004, Brown et al. 2008, and Kuhn et al. 2011).

including the nudge with savings of households that received an otherwise identical email without the nudge. A total of more than 40,000 emails were sent, of which a random half included the social norm nudge and the other half did not. Our experimental design allows us to establish causality.

There are two main mechanisms through which a social norm nudge can affect savings behavior: imitation behavior (Cialdini et al. 1990) and conformity preferences or identity considerations (Bernheim et al. 1994 and Akerlof and Kranton 2000). Imitation behavior predicts that households who receive the nudge increase their savings, because if most people have saved more, it must be “a sensible thing to do”. Conformity preferences and identity considerations give rise to an intrinsic disutility from not conforming to the social norm, for instance a disutility because of shame, guilt, or a feeling of not belonging to the group. Since we only target households who save less than the norm, both these mechanisms would predict that the social norm nudge increases savings.⁸

We obtain the following results. First, we find that those who receive the email including the nudge click more often on a link to their personal page where they can start or adjust an automatic savings plan. However, our analysis of detailed anonymized bank data spanning the period from the end of 2017 to September 2018 shows that the social norm nudge has no effect on actual savings, neither in the short run nor in the long run. Likewise, there is no discernible effect on the frequency or size of automatic savings transactions. Our null findings are quite precisely estimated.

We run a number of additional analyses and find very similar effects. For instance,

⁸For the nudge to have any effect, it is of course necessary that it provides new information to the recipients or that it makes the information more salient (Bordalo et al. 2013).

just like for the full sample, we find no effect for subsamples of households who arguably have more opportunities to save and for those who live in more homogeneous neighborhoods.⁹ Also, we find no difference in estimated treatment effects when we drop those clients from the sample who, at the moment of receipt of the email, have either a surprisingly high or low buffer.¹⁰ Results are also the same if we focus on a group of clients who open the email in the week that the emails were sent rather than later. Lastly, we do not find heterogeneous treatment effects across a set of demographics (gender, age, and household size).

One possible reason for our null finding on actual savings is a spillover effect of the treatment on households in the control group. Such spillovers may arise if control group households hear about and imitate saving plans of treated households. Households in the control group may also observe a change in consumption of treated households and, as a consequence, change their consumption as well (see Kuhn et al. 2011 for evidence of such peer effects in consumption expenditures in Dutch neighborhoods). By not taking these spillovers into account we may falsely conclude that the nudge has no effect, when in fact both the treatment group and the control group increase their savings due to the treatment. By virtue of our design, we can examine the existence of spillover effects. In the design stage, we followed the approach in Crépon et al. (2013) and created random variation in the fraction of households that receive treatment in each neighborhood. We find no evidence that saving behavior

⁹See Bicchieri and Dimant (2019) for a discussion of the importance of using a not-too-dissimilar reference group in social norm nudges.

¹⁰For technical reasons, the selection of households that was included in our field experiment took place a few weeks before the emails were sent. As a result, a small group of households ended up in our sample even though they no longer matched our selection criteria regarding the amount of their savings buffer. See Section 2 for more details.

of households in the control group varies with the fraction of households treated in their neighborhood, suggesting spillover effects are not important in our context.

Another possible reason for our null finding is that people simply did not notice the nudge. A complementary small-scale survey experiment conducted before the field experiment suggests otherwise, however. The survey experiment shows that the nudge attracts attention (measured using a software designed to track eye movements). Moreover, we find that the nudge increases savings intentions and encourages subjects to change their savings method. However, the nudge also leads to more annoyance.

Social norm nudges to increase household savings have been studied before by Beshears et al. (2015) and Kast et al. (2018). Kast et al. (2018) examine – in addition to several other interventions – the effect of repeatedly providing information about peers’ saving behavior to microcredit clients in Chile and find a sizeable effect on net new savings.¹¹ In contrast to our study, they cannot distinguish the impact of peer information from the effect of a pure reminder (see Karlan et al. 2016 and Rodríguez and Saavedra 2019 for evidence on the effect of reminders on savings). Our study is closest to Beshears et al. (2015) who examine the effect of a social norm nudge on retirement savings. In their field experiment, information about the retirement savings behavior of co-workers is added to plan enrollment and contribution increase forms that are distributed to low-saving workers. They find a negative effect of the intervention on retirement savings.¹² The reference group in

¹¹During the three month intervention period, net new savings are almost 80% higher in the treatment group than in the control group.

¹²Among the workers with a 0% contribution rate default, the intervention significantly decreased the average before-tax contribution rate from about 0.6% to about 0.4% during the month after

their intervention are co-workers of about the same age, but with sometimes very different wages. The negative effect of the nudge mainly arises for low-wage workers. Apparently, low-wage workers feel discouraged when being compared to colleagues who have higher economic status and for whom it is presumably easier to conform to the social norm of high retirement savings. Our paper revisits social norm nudging in the context of household finance. We try to avoid the discouragement effect found by Beshears et al. (2015) by selecting groups of households that are highly homogeneous. Yet, despite this, and despite the quite encouraging immediate proximate effects that we find, we find no effects on household savings, neither in the short run nor in the long run. While we were surprised by this result, our findings are in line with a number of recent studies that find small or negligible long-run welfare effects of nudges, despite sizeable short-run proximate effects, see among others Adams et al. (2018) in the context of credit card debt repayment, Choi et al. (2011) in the context of pension savings, and Loewenstein and Chater (2017) and Beshears and Kosowsky (2020) for a more general discussion.

Our paper proceeds as follows. In the next section we describe the design of the field experiment and discuss some descriptives. Section 3 describes our empirical strategy and reports the results of the field experiment. Section 4 describes the set-up and reports the results of our complementary survey experiment. Section 5 concludes.

the treatment.

2 Experimental Design

The randomized field experiment took place at ING Netherlands, a large retail bank that has close to 8 million clients in the Netherlands (that is, nearly half of the Dutch population). We set up an email marketing campaign encouraging households with a low savings buffer to save more. More specifically, the campaign targeted households that have a savings buffer that is smaller than the median savings buffer of other ING clients living in the same neighborhood, where savings buffer is defined as the sum of the amounts on the current and liquid savings accounts at ING. Households in the treatment group receive an email containing a social norm nudge, those in the control group receive an otherwise identical email without the nudge. The social norm nudge is displayed in a picture in the email and reads: “You have a lower buffer with us than most other ING clients in your neighborhood”. Figure 1 shows the control email (on the left-hand side) and the treatment email (on the right-hand side).

The texts below the pictures in Figure 1 are identical and read:

“Dear [Mr. or Ms. Last Name],

Although the interest rate is low, saving offers the certainty of a buffer that you can always use.

Save automatically

Did you know that you can save almost effortlessly? Set automatic saving once, you can do that in two minutes. You choose the amount, the frequency, and the end date. Done.

Do it directly. With automatic saving, you build a buffer effortlessly.

Set it now >

The final sentence "Set it now >" is displayed in an orange box and links to a personal webpage where the client can start or adjust an automatic savings plan. The grey bar below the orange box contains two links, one to opt out of any future marketing emails from ING ("> Afmelden") and the other to notify ING of a change in email address ("> E-mail gegevens wijzigen"). For all three links it holds that the client needs to log in with a username and password to access the personalized website. Lastly, the subject lines of the emails were also identical in treatment and control and read "What if things go wrong?".

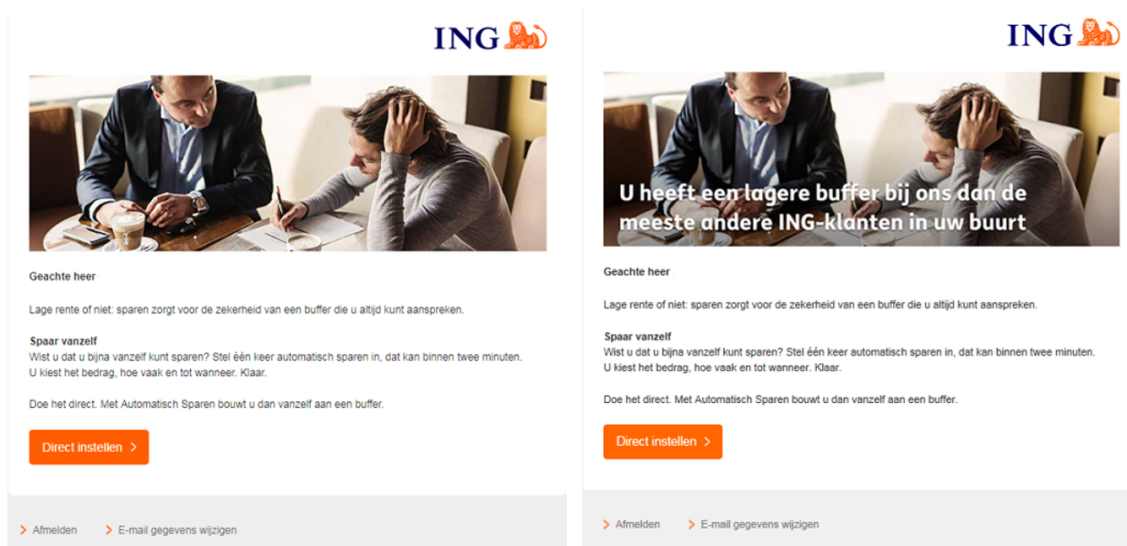


Figure 1: The email sent to households in the control group (on the left-hand side) and the treatment group (on the right-hand side). Note: Full translation of the emails is provided in the main text.

The social norm nudge in the treatment email makes a comparison of the client’s buffer savings with buffer savings of other ING clients in the same neighborhood. We define neighborhoods as five-digit zip code areas, implying that in each neighborhood on average about 250 households reside (of which about half are ING clients). The Netherlands consists of 33,000 of such neighborhoods.

ING Data Protection Board allowed us to analyze detailed anonymized microdata for a maximum of 15,000 clients. The sample we analyze consists of households that opened the email and loaded the picture that is in the email, which is tracked by ING. Since the treatment message was included in the picture, and all other parts of the emails (including the subject line) were identical between treatment and control, there can be no selection into treatment.

The bank knows from previous promotional campaigns that about 40% of the clients open marketing emails from ING. Therefore, to achieve a final sample of 15,000 clients, we selected slightly more than 40,000 clients for the email marketing campaign and did so in the following way. Firstly, we selected households that have a savings buffer that is smaller than the median savings buffer of ING clients living in the same neighborhood. Secondly, we excluded households that seem to have little opportunity to increase their savings. More specifically, we imposed the restriction that a household should have a regular inflow of money of at least 1,000 euros per month, which is close to the minimum welfare benefit for a single person in the Netherlands. This requirement also makes it likely that for the clients we selected the ING account is the primary bank account.¹³ Thirdly, we excluded households

¹³Clearly, our social norm nudge is less relevant for clients who do not only have a bank account with ING, but also with other banks. According to the household survey of the Dutch central bank

with a negative savings buffer. Fourthly, we dropped all households for which no email address is available, who have opted out of receiving any marketing emails, or who have recently received another marketing email from ING Netherlands. Lastly, to minimize the discouragement effect arising from upward social comparison that was found in Beshears et al. (2015), we focused on households who live in relatively homogeneous neighborhoods. We selected 1,904 neighborhoods where households are most similar in terms of the type of houses, the age distribution, and a proxy for the income distribution. Online Appendix A.1 describes in detail how we did this. The selection resulted in a sample of 41,602 households, who were all sent either a control or treatment email in the first week of 2018.¹⁴

Randomization of households into treatment and control was done following the approach set out in Crépon et al. (2013), which allows to detect possible contamination of the control group. We expect possible spillovers to be the largest within the neighborhood. Therefore, we varied treatment intensity by neighborhood. In a random half of the neighborhoods, we randomly assigned 80% to treatment and 20% to control. In the other half, we randomly assigned 20% to treatment and 80% to control. Our randomization thus took place at two levels. First we randomized at the neighborhood level (randomizing over the 1,904 neighborhood), and then at the households level (within each neighborhood). If control households' savings behavior is related to the share of treated households in their neighborhood, we take this as an indication for contamination.

(DNB Household Survey), a majority of ING clients do not hold a current account at another bank.

¹⁴Slightly more than 15,000 clients opened the email, so we needed to make a further selection to meet the restriction of the ING Data Protection Board. Online Appendix A.1 describes how we did this.

Note that we did not create a control group of clients who receive no email at all. As a result, we can only estimate the effect of the social norm nudge, while we can not estimate the effect of receiving the email as compared to not receiving the email. While this would have been an interesting addition to our study, the restriction on the total sample imposed by ING Data Protection Board implies that it would have reduced the power of our nudge experiment. Using a significance level of 0.05 and setting power to 0.80, our experiment’s minimum detectable effect size on buffer savings is 181 euros, which is equivalent to 0.05 of a standard deviation and 0.09 of the control mean.¹⁵

Our key variable of interest is the households’ buffer savings, for which we have weekly data starting 19 weeks before the intervention up to and including 38 weeks after the intervention (the amounts on Sunday each week). Table 1 provides some descriptive statistics of the savings data on Sunday 31 December 2017, a few days before the intervention. Households in our sample have on average a buffer of about 2,000 euros (the median is about 1,400 euros).¹⁶ In addition to the effect on buffer savings, we also examine the effect of the social norm nudge on whether people make an automatic savings transaction and on the amount saved automatically. The week before the intervention, about one sixth of households made an automatic savings transaction. On average, households automatically transferred slightly more than

¹⁵The standard deviation of buffer savings is quite large as a result of a small number of outliers. On the last Sunday before the intervention, the standard deviation of buffer savings is 3,964 euros. The standard deviation is much lower when dropping the top and bottom 1% of the observations. In that truncated sample, the standard deviation is 1,680 euros.

¹⁶The standard deviations are large and differ quite a bit between treatment and control as a result of outliers. When we drop the top and bottom 1% of the observations, the standard deviation is 1,675 euros and 1,686 euros in treatment and control, respectively.

twenty euros to their savings account. None of these savings data differ significantly between the treatment and control group. The same holds for our demographic variables household size, gender, and age, which are also shown in Table 1. On average, a household has 2.1 members, 55% is male, and the average age of the account holder is close to 47 years.¹⁷

Table 1: Descriptive statistics of savings data and demographics on the last Sunday before treatment

	Control $N=7,458$	Treatment $N=7,542$	p-value
Buffer savings	2015.86 (2922.95)	2042.87 (4774.85)	0.68
Automatic savings transaction	0.167 (0.373)	0.162 (0.369)	0.39
Automatic savings amount	23.19 (84.70)	24.26 (93.34)	0.46
Household size	2.12 (1.13)	2.12 (1.14)	0.89
Gender (male = 1)	0.56 (0.50)	0.55 (0.50)	0.74
Age	46.81 (14.33)	46.48 (14.58)	0.16

Standard deviations are in parentheses, p-values < 0.10, 0.05, 0.01 respresented by *, **, ***, respectively. Buffer savings is the sum of amounts on all current and savings accounts at ING of a household in euros. Automatic savings transaction is a dummy equal to 1 if at least one automatic savings transaction at ING took place during the last 7 days. Automatic savings amount is the amount that has been automatically transfered to a savings account at ING in the last 7 days.

¹⁷We only observe household members if they have an account at ING, and so household size is in fact larger for some households. The age and gender is of the person who received the email. If there is more than one ING client in a household, the email is sent to the person who is registered as the main contact — and if none is registered as main contact, to a randomly selected adult household member.

3 Results

This section describes the key results of our field experiment. We start with the immediate proximate effects, i.e. clicks on the links in the email. Next we study the effects on household savings in the short, medium, and long run. Finally, we examine heterogeneous treatment effects and spillovers.

3.1 Immediate proximate effects

Table 2 shows the immediate impact of the social norm nudge as measured by clicks on the links in the emails during the first week after sending them. We consider the clicks on two links: the link to go to a personalized website to start or adjust automatic savings and the link to opt out of future emails. We find that a significantly larger share of households in the treatment group clicked on the link to start or adjust automatic savings than in the control group (3.4% versus 2.7%). We find no difference in clicks on the link to opt out of future marketing email messages from ING (0.5% opts out in both the treatment and the control group). These findings suggest that the social norm nudge had the intended effect of encouraging households to start or adjust automatic savings, while at the same time engendering little annoyance. The latter finding is in contrast to the recent study by Damgaard and Gravert (2018) who find that nudging increases unsubscriptions from a mailing list in a field experiment with a charity.

Table 2: Immediate proximate effects

	Control $N=7,458$	Treatment $N=7,542$	p-value
Click to automatic savings page	0.027 (0.162)	0.034 (0.180)	0.00***
Click to opt out of future emails	0.005 (0.069)	0.005 (0.069)	0.99

Standard deviations are in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively. The click to automatic savings page and the click to opt out of future emails are dummy variables.

3.2 Effects on savings

Our primary outcome variable is the amount of buffer savings. Before turning to regression analysis, we first provide some graphical evidence.

Figure 2 plots average buffer savings of households in the treatment and control group separately during our full sample period. The black vertical line indicates the timing of the intervention. It is clear that the savings buffer varies substantially by week for both the treatment and control group. This can be explained by recurrent monthly events such as receipt of salary or benefit as well as monthly expenditures such as rent, mortgage, and insurance premiums.¹⁸ Further, by comparing the two lines, we see before as well as after the intervention only minor differences in the averages for control and treatment group households. This indicates that the randomization was successful in balancing the two groups, but also that there does not seem to be an average treatment effect of the social norm nudge, neither in the short run nor in the long run. This becomes even more clear in Figure 3, which plots

¹⁸The peak in Figure 2 at the end of the year (two weeks before treatment) is probably largely due to year-end bonuses, while the peak in May (the fifth spike after treatment) can be attributed to holiday pay.

the difference in average buffer savings between treatment and control group over time. If anything, there seems to be a small negative treatment effect during the first twenty weeks of about 50 euros, which is equivalent to about 0.02 of a standard deviation.




Figure 2: Average buffer savings in euros of households in the treatment and control group. Note: the black vertical line denotes the time the email was sent.




Figure 3: Difference in average buffer savings in euros between households in the treatment and those in the control group. Note: the black vertical line denotes the time the email was sent.

Table 3 reports the estimation results of regressing the amount of buffer savings on a treatment dummy, the amount of buffer savings the week before the intervention, and a constant. We run the regression for buffer savings at three points in time: four weeks, 13 weeks, and 38 weeks after the intervention (which is the last week for which we have data).¹⁹ In line with the graphical evidence shown above, we find no significant treatment effect, neither in the short run nor in the long(er) run.²⁰

¹⁹Thirteen clients (less than 0.1% of the total sample) no longer shows up in the data at some point after treatment because their account has been closed. We checked whether there is an effect of the treatment on attrition, but find none ($p=0.78$).

²⁰The standard errors show that our estimates are quite precise, but precision does decline over time. We greatly gain in precision when truncating the data: dropping the observations that belong to the top or bottom 1% of the distribution of savings results in a standard error of the coefficient for the treatment dummy of about 25 for the regression of buffer savings at $t=4$, 34 at $t=13$, and 52 at $t=38$.

Table 3: Effect of the social norm nudge on buffer savings

	Buffer savings at t=4	Buffer savings at t=13	Buffer savings at t=38
Treatment	-65.36 (63.96)	-59.91 (82.85)	24.47 (132.80)
Buffer savings at t=0	0.90*** (0.01)	0.42*** (0.01)	0.78*** (0.02)
Constant	805.49*** (48.14)	1592.76*** (62.41)	1181.01*** (100.10)
Number of households	15,000	14,999	14,987
R ²	0.45	0.10	0.13

Standard errors are in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

The mean and standard deviation of buffer savings at t=0 are equal to 2029.44 euro and 3963.65 euro, respectively.

In Tables 4 and 5, we follow the same approach to study the effects of the social norm nudge on the use and amount of automatic savings and find statistically insignificant effects that are very close to zero, both in the short run and in the long(er) run.

Table 4: Effect of the social norm nudge on the amount put in a savings account

	Auto savings amount at t=4	Auto savings amount at t=13	Auto savings amount at t=38
Treatment	-0.76 (1.88)	-0.75 (1.54)	-1.85 (1.17)
Auto savings amount at t=0	1.16*** (0.01)	1.19*** (0.01)	0.04*** (0.01)
Constant	10.75*** (1.35)	8.94*** (1.11)	11.09*** (0.84)
Number of households	15,000	14,999	14,987
R ²	0.45	0.56	0.01

Standard errors are in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The variable

"auto savings amount" denotes the amounts on the savings account in euros.

stock at 0: mean: 23.73, standard deviation 89.15.

Table 5: Effect of the social norm nudge on making an automatic savings transaction

	Automatic transaction at t=4	Automatic transaction at t=13	Automatic transaction at t=38
Treatment	-0.006 (0.005)	-0.002 (0.004)	-0.007* (0.004)
Automatic transaction at t=0	0.604*** (0.006)	0.779*** (0.005)	0.073*** (0.005)
Constant	0.055*** (0.003)	0.035*** (0.003)	0.055*** (0.003)
Number of households	15,000	14,999	14,987
R ²	0.39	0.61	0.01

Standard errors are in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

Automatic transaction is a dummy variable that equals 1 in case at least one automatic savings transaction took place in the preceding week.

The mean and standard deviation of automatic transaction at t=0 are equal to 0.165 and 0.371, respectively.

3.3 Heterogeneous Treatment Effects

The zero average treatment effects we found may hide interesting heterogeneous treatment effects. In Tables A1-A5 of the Online Appendix, we examine whether the treatment effects differ for several subsamples and find that they do not. We report treatment effects 13 weeks after the treatment for the following subsamples: i) Those who open the email in the first five days after sending it (96% of the sample), as these subjects may have a stronger interest in emails from ING and may pay more attention to it (Table A1); ii) Those for whom our selection criteria apply more broadly, namely up to and including two months before the treatment (89% of the sample, Table A2); iii) Those living in the top half most homogeneous neighborhoods (53% of the sample), as the social norm nudge may be considered as more relevant by subjects whose neighbors are more similar to them (Table A3); iv) Those who seem to have more opportunities to increase savings as proxied by the difference between the

total inflow of money and the automatic transfers (e.g. rent and energy payments) in the month before the intervention (various subsamples, see Table A4); and v) Those with a non-negative buffer three days before the email arrives (97% of the sample), and those whose buffer is less than 10,000 euros three days before the email arrives (99.9% of the sample), see Table A5. While the lack of heterogeneous treatment effects may seem surprising at first sight, note that it is at least partly implied by our design. As explained in section 2, we deliberately selected households living in highly homogeneous neighborhoods, who have little buffer savings, and who seem to have some opportunities to save more. Our analysis of heterogeneous treatment effects exploits the remaining variation in the selected sample, which is limited.

Tables A6-A10 of the Online Appendix study heterogeneous treatment effects across several demographics, including household size (Table A6), gender (Tables A7 and A8), and age (Tables A9 and A10). For none of these subsamples do we find treatment effects that are statistically significantly different from zero at conventional levels.

3.4 Spillover Effects

One reason for why the estimated effect of the social norm nudge are so close to zero is contamination of the control group. By virtue of our design, we can shed light on whether control group subjects have been affected by the treatment. Following the approach in Crépon et al. (2013) we varied the fraction of households who receive treatment by neighborhood. More specifically, in some randomly chosen neighborhoods (which we call “high-dose treatment”) 80% were randomly assigned

to the treatment group, while the remaining 20% were randomly assigned to the control group. In the other randomly chosen neighborhoods (which we call “low-dose treatment”), 20% were randomly assigned to treatment, while the remaining 80% were randomly assigned to the control group. If contamination is important, we expect that savings behavior of control households is related to the fraction of treated households in their neighborhood.

To explore whether that is the case, we estimate exactly the same regression as in the previous subsections, except that we replace the treatment dummy by three interaction terms: the treatment dummy interacted with a dummy for “high-dose treatment”, the treatment dummy interacted with a dummy for “low-dose treatment”, and finally a dummy for the control group interacted with a dummy for “high-dose treatment”. If contamination is important, we expect the coefficient on this latter interaction term to be positive, i.e. savings in control neighborhoods where a high fraction of households were treated are higher than savings in control neighborhoods where a low fraction of households were treated.

Table 6 shows the results for our three key measures of savings behavior, all measured three months after the intervention. We find no statistically significant spillover effects on the control group, see the estimates in the first row of the table.

Table 6: Spillover effects of the social norm nudge 3 months after treatment

	Buffer savings	Automatic savings amount	Automatic savings transaction
Control × High-dose treatment	-188.58 (145.93)	-4.79* (2.71)	-0.011 (0.007)
Treatment × Low-dose treatment	63.90 (146.27)	1.96 (2.72)	-0.001 (0.002)
Treatment × High-dose treatment	-138.79 (92.72)	-2.65 (1.72)	-0.005 (0.004)
Buffer savings at t=0	0.42*** (0.01)	1.19*** (0.01)	0.779*** (0.005)
Constant	1631.55*** (69.22)	9.92*** (1.24)	0.037*** (0.003)
Number of households	14,999	14,999	14,999
R ²	0.10	0.56	0.61

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

4 Survey Experiment

The large-scale field experiment was preceded by a small-scale survey experiment, used as a pilot, to study the following three issues: i) Does the nudge attract the clients' attention?; ii) Do clients increase their savings intentions in response to the nudge?; and iii) Is the nudge perceived as intrusive by clients?²¹ 292 people participated in the survey, of which 147 saw the treatment email and 145 saw the control email, see Figure 1. The survey was developed and administered by an external research bureau (DVJ Insights) in cooperation with a communication researcher of ING, Yoka Wesseling. The participants were selected from an existing panel run by the bureau. The bureau selected participants who are similar to the people in our

²¹The survey experiment examined, in addition to our treatment email and control email, two other emails with other versions of the social norm. One of these provided more detailed information about how the client's buffer savings compares to that of people in the neighborhood, while the other email provided coarser information. Both of these alternative emails were judged as less motivating as compared to the treatment email we sent in the field experiment. In the following, we ignore these two other conditions of the survey experiment for the sake of brevity.

field experiment: clients at ING, with low savings, sufficient income, and living in homogeneous neighborhoods. Survey participants were randomly selected to see one of the messages and were subsequently asked questions about their motivation to save, the intrusiveness of the message, as well as his or her perceptions about the bank.

In order to test whether the nudge is actually noticed, a software developed by DVJ Insights was used that instructs the respondent to move the cursor to the position on the screen where one is looking. Figure A.1 in Online Appendix A.3 shows a heatmap suggesting that respondents in the treatment group spend more time looking at the picture, and particularly so at the part that displays the social norm nudge. The data in Figure A.2 in Online Appendix A.3 suggest that approximately 76% of the respondents in the treatment group reads the social norm nudge. This comes at the cost of attention paid to the main text in the email: This is read by 81% in the control group versus 69% in treatment. This might be due to respondents having limited time or attention span – implying that adding information at the top part of the email crowds out attention paid to the bottom – or because the social norm nudge reveals information about the likely content of the email, and so may discourage some people to read the full email. Despite the drop in readership of the body text, a slightly *higher* percentage of respondents in treatment looks at the very bottom of the email where the link to the automatic savings page is. The social norm nudge apparently induced respondents to pay attention to the part of the email where they can follow up on the message they have just read. A possible reason is that the social norm nudge increased the relevance of the subject to these respondents.

Table 7 provides some support for this interpretation. It shows the mean of answers given to a range of questions of respondents in the treatment and control group. Clearly, respondents in the treatment group find the email more relevant (22% versus 13%). Moreover, the social norm nudge increases respondents' intention to save. 19% of respondents in treatment is motivated to save more versus 12% in control. Also, a higher percentage of people states to be prepared to save automatically each month (22% versus 14%) and to change the current saving method (18% versus 8%). While this evidence is quite encouraging, Table 7 also makes clear that the motivating power of the social norm nudge does not come for free. A higher percentage of respondents in the treatment group is annoyed by the email (33% versus 13%) and a higher percentage finds the email unacceptable (26% versus 9%).²² Despite this, there do not seem to be major repercussions for the image of the bank, see the bottom part of Table 7.²³

²²In open response boxes, several respondents write that they are annoyed by the email because they do not want the bank to compare them with people in their neighborhood and because they find it unacceptable that the bank shares information about buffer savings of people in their neighborhood with them (even though it is coarse and aggregated). Interestingly, in the treatment group, there is a strong negative correlation (-0.16) between whether people find the email annoying and whether it motivates them to save more. Likewise, those who find the email motivating, also find the email less often unacceptable (the correlation is -0.21). In the control group, these correlations are much weaker and statistically insignificant (-0.08 and -0.11, respectively).

²³Note that a majority of the sample considers the bank as reliable, and this does not differ between control group and treatment group. Bicchieri and Dimant (2019) argue that for social norm nudges to work the sender needs to be a trusted source of information. Consistent with this, in our survey data we find that those who find the bank reliable are more likely to report that the email motivates them to save more (the correlation is 0.18 in the treatment group and 0.15 in the control group).

Table 7 Survey experiment: participants' responses to the survey questions by treatment

	Control <i>N</i> =145	Treatment <i>N</i> =147	p-value
This message:			
motivates me to save more	0.12	0.19	0.08*
I am prepared to save automatically each month	0.14	0.22	0.08*
induces me to change my current saving method	0.08	0.18	0.01***
This message:			
is positive	0.26	0.20	0.24
is relevant	0.13	0.22	0.05**
is credible	0.28	0.28	0.94
is understandable	0.55	0.60	0.35
is annoying	0.13	0.33	0.00***
is informative	0.27	0.29	0.66
is unacceptable	0.09	0.26	0.00***
This is a bank that:			
prioritises the interest of the clients	0.32	0.35	0.59
is reliable	0.57	0.53	0.55
does what she promises	0.45	0.44	0.83
is open and transparent about how she operates	0.29	0.37	0.16
has integrity	0.48	0.46	0.65
is careful with my personal information	0.57	0.53	0.55

p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively. Columns 1 and 2 give the mean proportion (strongly) agreeing with the statement (i.e. score 6 or higher on a 7-point-scale).

5 Conclusion

We ran a large-scale field experiment at a retail bank in the Netherlands to study the impact of a social norm nudge on buffer savings by households. We added a social norm nudge to an email message from the bank promoting savings. We found that adding the nudge increases households' savings intentions as measured by the responses of subjects in a survey experiment. Moreover, using data on clicks and website visits from the field experiment, we found that adding the nudge induces households to take some steps towards changing their savings method or savings

amount. However, using detailed anonymized bank data, we found no effect on either the amount of buffer savings or on the way people save, neither in the short run nor in the long run. These results are surprising given the existing body of evidence on peer effects in household financial decision making and given the successes that have been achieved with social norm nudging in changing people’s behavior in other contexts that we discussed in the Introduction. The results are even more surprising given that we have selected a sample that seems very fertile ground to find a social norm nudge effect. While it is hard to point to the exact reasons for the lack of an effect, we ruled out a number of candidate explanations such as a lack of attention to the nudge, insufficient opportunities for households to adjust their savings, or a massive lack of trust in the sender of the message.²⁴

In addition to these contributions, our study can also serve as a reminder that intentions as expressed in surveys (or revealed by clicks on links to a webpage) are not always followed up by substantive change in behavior. Choi et al. (2006) make a similar point in a related context. They found that many participants to a financial education seminar stated that they would change their retirement savings behavior, but most of them did not actually change anything after all. Using data other than actual decisions may thus lead researchers to falsely claim that a treatment has an effect.

Even though we were surprised that we did not find a positive effect of the

²⁴In a recent related study, Goldin et al. (2020) stress the role of decision costs in saving choices and provide evidence from a field experiment on retirement saving in support of this idea. In the spirit of their study, we aim to provide personalized advice on how much to save in future field experiments. In such a follow-up study, it would also be interesting to allow for a pure control group that does not receive any message, so as to study the effects of reminding people to save, as in Karlan et al. (2016) and Rodríguez and Saavedra (2019).

social norm nudge on savings, our findings are consistent with a number of recent studies finding small or negligible welfare effects of nudges in the long run, even when proximate outcomes in the short run sometimes look very promising. In addition to the studies mentioned in the penultimate paragraph of the Introduction, DellaVigna and Linos (2021) have recently shown that the effects of nudges in trials performed by Nudge Units in governments are much smaller than the effects of nudges reported in the academic literature. While this evidence is sobering, we should keep in mind that the costs of nudges are typically also very low, making the use of nudges oftentimes socially optimal, even if they have only small effects (Benartzi et al. 2017; Laibson 2020). For the same reason, it is clear that we should not give up on testing nudges using randomized controlled trials, particularly when highly important issues such as improving financial resilience of households are at stake.

Acknowledgements

We are most grateful to Marieke Blom, Jesse Pouw, Yoka Wesseling, DVJ Insights, and the Think Forward Initiative. Without them we would not have been able to perform this study. We thank a co-editor and three reviewers of this journal for extremely useful comments on a previous draft. We also gratefully acknowledge comments and suggestions by Paul Adams, Jim Been, Eugen Dimant, Sacha Kapoor, Stephan Meier, Paul Smeets, Eduard Suari Andreu, Benedikt Vogt, and workshop and conference participants of the Field Days 2018 in Rotterdam, the International Pension Workshop 2019 in Leiden, the Maastricht Behavioral and Experimental Economics Symposium 2019, the WINK Conference 2019 in Utrecht, the Natural Experiments and Controlled Field Studies Workshop 2019 in Ohlstadt, the KVS New Paper Sessions 2019 in The Hague, and seminar participants at the Netherlands Authority for Financial Markets, CPB Netherlands Bureau for Economic Policy Analysis, the Shandong University of Finance and Economics in Beijing, Fudan University in Shanghai, and the Central University of Finance and Economics in Jinan.

References

- [1] Adams, P., Guttman-Kenney, B., Hayes, L., Hunt, S., Laibson, D., & Stewart, N. (2018). The semblance of success in nudging consumers to pay down credit card debt. FCA Occasional Paper No. 45.
- [2] Akerlof, G. A., & Kranton, R. E. (2000). Economics and Identity. *Quarterly Journal of Economics*, 115(3), 715-753.
- [3] Allcott, H. (2011) Social Norms and Energy Conservation. *Journal of Public Economics*, 95(9-10), 1082–1095.
- [4] Allcott, H., & Rogers, T. (2014). The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation. *American Economic Review*, 104(10), 3003-3037.
- [5] Bell, C., Gorin, D., & Hogarth, J. M. (2008). "Financial Education - Does It Work and How Do We Know? Research Findings from a Study of Financial Education Among Soldiers." *Community Invest* 21, 15–16.
- [6] Benartzi, S., Beshears, J., Milkman, K.L., Sunstein, C.R., Thaler, R.H., Shankar, M., Tucker-Ray, W., Congdon, W.J., Galing, S. (2017). Should Governments Invest More in Nudging? *Psychological Science*, 28(8), 10411055.
- [7] Bernheim, B. D. (1994). A Theory of Conformity. *Journal of Political Economy*, 102(5), 841-877.

- [8] Bernheim, B. D., & Garrett, D. (2003). "The Effects of Financial Education in the Workplace: Evidence from a Survey of Households." *Journal of Public Economics*, 87(7–8), 1487–1519.
- [9] Beshears, J., Choi, J.J., Laibson, D., Madrian, B.C., & Milkman, K.L. (2015). "The Effect of Providing Peer Information on Retirement Savings Decisions." *Journal of Finance*, 70(3), 1161-1201.
- [10] Beshears, J., Choi, J.J., Laibson, D., & Madrian, B.C. (2018). "Behavioral Household Finance." In: Bernheim, D.B., DellaVigna, S., & Laibson, D. (Eds.), *Handbook of Behavioral Economics*, 177-276, Amsterdam: Elsevier.
- [11] Beshears, J., & Kosowsky, H. (2020). Nudging: Progress to Date and Future Directions. *Organizational Behavior and Human Decision Processes*, 161(Supplement): 319.
- [12] Bhanot, S.P. (2018). Isolating the Effect of Injunctive Norms on Conservation Behavior: New Evidence from a Field Experiment in California. *Organizational Behavior and Human Decision Processes*, forthcoming.
- [13] Bhutta, N., & Dettling, L. (2018). Money in the Bank? Assessing Families' Liquid Savings using the Survey of Consumer Finances. FEDS Notes, Board of Governors of the Federal Reserve System.
- [14] Bicchieri, C., Dimant, E. (2019). Nudging with Care: The Risks and Benefits of Social Information. *Public Choice*, forthcoming.

- [15] Blumenthal, M., Christian, C., & Slemrod, J. (2001). Do Normative Appeals Affect Tax Compliance? Evidence from a Controlled Experiment in Minnesota. *National Tax Journal*, 54(1), 125-138.
- [16] Bordalo, P., Gennaioli, N., Shleifer, A. (2013). Salience and Consumer Choice. *Journal of Political Economy*, 121(5), 803-843.
- [17] Bott, K. M., Cappelen, A. W., Sorensen, E., & Tungodden, B. (2020). You've Got Mail: A Randomised Field Experiment on Tax Evasion. *Management Science*, 66(7): 2801-2819.
- [18] Bradler, C., Dur, R., Neckermann, S., & Non, A. (2016). Employee Recognition and Performance: A Field Experiment. *Management Science*, 62(11), 3085-3099.
- [19] Brandon, A., Ferraro, P. J., List, J. A., Metcalfe, R. D., Price, M. K., & Rundhammer, F. (2017). Do the Effects of Social Nudges Persist? Theory and Evidence from 38 Natural Field Experiments. NBER Working Paper No. w23277.
- [20] Brown, J.R., Farrell, A.M., & Weisbenner, S.J. (2016). Decision-Making Approaches and the Propensity to Default: Evidence and Implications. *Journal of Financial Economics*, 121(3), 477-495.
- [21] Brown, J. R., Ivkovi, Z., Smith, P. A., Weisbenner, S. (2008). Neighbors Matter: Causal Community Effects and Stock Market Participation. *Journal of Finance*, 63(3), 1509-1531.

- [22] Brown, K. M., Laschever, R. A. (2012). When They're Sixty-Four: Peer Effects and the Timing of Retirement. *American Economic Journal: Applied Economics*, 4(3), 90-115.
- [23] Bursztyn, L., Ederer, F., Ferman, B., and Yuchtman, N. (2014). Understanding Mechanisms Underlying Peer Effects: Evidence from a Field Experiment on Financial Decisions. *Econometrica*, 82(4), 1273-1301.
- [24] Chen, Y., Harper, F.M., Konstan, J., & Li, S.X. (2010). Social Comparisons and Contributions to Online Communities: A Field Experiment on Movielens. *American Economic Review*, 100(4), 1358-1398.
- [25] Choi, J, Laibson, D., Madrian B., and Metrick A. (2006). "Saving for Retirement on the Path of Least Resistance." In: McCaffrey, E., and Slemrod, J. (Eds.), *Behavioral Public Finance: Toward a New Agenda*, 304-351. New York: Russell Sage Foundation.
- [26] Choi J., Laibson, D., & Madrian, B. (2011). \$100 Bills on the Sidewalk: Sub-optimal Investment in 401(k) Plans. *Review of Economics and Statistics*, 93(3), 748-763.
- [27] Cialdini, R.B., Reno, R.R., & Kallgren, C.A. (1990). A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places. *Journal of Personality and Social Psychology*, 58(6), 1015.

- [28] Coffman, L.C., Featherstone, C.R., & Kessler, J.B. (2017). Can Social Information Affect What Job You Choose and Keep? *American Economic Journal: Applied Economics*, 9(1), 96-117.
- [29] Cranor, T., Goldin, J., Homonoff, T., Moore, & L. (2020). Communicating Tax Penalties to Delinquent Taxpayers: Evidence from a Field Experiment. *National Tax Journal*, 73(2), 331-285.
- [30] Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., & Zamora, P. (2013). Do Labor Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment. *Quarterly Journal of Economics*, 128(2), 531-580.
- [31] Damgaard, M.T., & Gravert, C. (2018). The Hidden Costs of Nudging: Experimental Evidence from Reminders in Fundraising. *Journal of Public Economics*, 157, 15-26.
- [32] DellaVigna, S., & Linos, E. (2021). RCTs to Scale: Comprehensive Evidence from Two Nudge Units. Mimeo, University of California, Berkeley.
- [33] Dimant, E., Van Kleef, G.A., & Shalvi, S. (2020). Requiem for a Nudge: Framing Effects in Nudging Honesty. *Journal of Economic Behavior & Organization*, 172, 247-266.
- [34] Duflo, E., & Saez, E. (2002). Participation and Investment Decisions in a Retirement Plan: The Influence of Colleagues Choices. *Journal of Public Economics*, 85(1), 121-148.

- [35] Duflo, E., & Saez, E. (2003). The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment. *Quarterly Journal of Economics*, 118(3), 815-842.
- [36] Fellner, G., Sausgruber, R., & Traxler, C. (2013). Testing Enforcement Strategies in the Field: Threat, Moral Appeal and Social Information. *Journal of the European Economic Association*, 11(3), 634-660.
- [37] Frey, B.S., & Meier, S. (2004). Social Comparisons and Pro-social Behavior: Testing "Conditional Cooperation" in a Field Experiment. *American Economic Review*, 94(5), 1717-1722.
- [38] Giaccherini, M., Herberich, D.H., Jimenez-Gomez, D., List, J.A., Ponti, G., & Price, M.K. (2019). The Behavioralist Goes Door-To-Door: Understanding Household Technological Diffusion Using a Theory-Driven Natural Field Experiment. NBER Working Paper No. 26173.
- [39] Gerber, A.S., & Rogers, T. (2009). Descriptive Social Norms and Motivation to Vote: Everybody's Voting and So Should You. *Journal of Politics*, 71(1), 178-191.
- [40] Goldin, J., Homonoff, T., Patterson, R., Skimmyhorn, W. (2020). How much to save? Decision costs and retirement plan participation. *Journal of Public Economics*, 191, 104247.

- [41] Goldstein, N., Cialdini, R., & Griskevicius, V. (2008). A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels. *Journal of Consumer Research*, 35(3), 472–482.
- [42] Hallsworth, M., Chadborn, T., Sallis, A., Sanders, M., Berry, D., Greaves, F., Clements, L., Davies, S.C. (2016). Provision of Social Norm Feedback to High Prescribers of Antibiotics in General Practice: A Pragmatic National Randomised Controlled Trial. *The Lancet*, 387(10029), 1743-1752.
- [43] Hallsworth, M., List, J.A., Metcalfe, R.D., & Vlaev, I. (2017). The Behavioralist as Tax Collector: Using Natural Field Experiments to Enhance Tax Compliance. *Journal of Public Economics*, 148, 14–31.
- [44] Hong, H., Kubik, J.D., & Stein, J.C. (2004). Social Interaction and Stock-Market Participation. *Journal of Finance*, 59(1), 137-163.
- [45] John, P., & Blume, T. (2018). How Best to Nudge Taxpayers? The Impact of Message Simplification and Descriptive Social Norms on Payment Rates in a Central London Local Authority. *Journal of Behavioral Public Administration*, 1(1), 1-11.
- [46] Kuhn, P., Kooreman, P., Soetevent, A., & Kapteyn, A. (2011). The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery. *American Economic Review*, 101(5), 2226–2247.

- [47] Karlan, D., McConnell, M., Mullainathan, S., Zinman, J. (2016). Getting to the Top of Mind: How Reminders Increase Saving. *Management Science*, 62(12), 3393-3411.
- [48] Kast, F., Meier, S., & Pomeranz, D. (2018). Saving More in Groups: Field Experimental Evidence from Chile. *Journal of Development Economics*, 133, 275-294.
- [49] Laibson, D. (2020). Nudges Are Not Enough. AEA Webcast, <https://www.aeaweb.org/webcasts/2020/aea-afa-joint-luncheon-nudges-are-not-enough>.
- [50] Lieber, E.M.J., Skimmyhorn, W. (2018). Peer Effects in Financial Decision-Making. *Journal of Public Economics*, 163, 37-59.
- [51] Loewenstein, G., & Chater, N. (2017). Putting nudges in perspective. *Behavioural Public Policy*, 1(1): 26-53.
- [52] Lusardi, A. (2004). "Saving and the Effectiveness of Financial Education." In: Mitchell, O.S, & Utkus, S.P. (Eds.), *Pension Design and Structure: New Lessons from Behavioral Finance*, 157-184, Oxford University Press, New York, NY.
- [53] Nibud (2017) 2,5 miljoen huishoudens hebben te weinig geld achter de hand [2,5 million households have too little buffer savings]. Online publication: <https://www.nibud.nl/beroepsmatig/nibud-25-miljoen-huishoudens-geld-hand/>

- [54] Ouimet, P., & Tate, G. (2020). Learning from Coworkers: Peer Effects on Individual Investment Decisions. *Journal of Finance*, 75(1): 133-172.
- [55] Rodríguez, C., & Saavedra, J.E. (2019). The Persistent Effects of Youth Savings Reminders: Experimental Evidence from Text-Message Campaigns in Colombia, *Journal of Development Economics* 139, 135-156.
- [56] Schultz, P.W., Nolan, J.M., Cialdini, R.B., Goldstein, N.J., & Griskevicius, V. (2007). The Constructive, Destructive, and Reconstructive Power of Social Norms. *Psychological Science*, 18(5), 429-434.
- [57] Skimmyhorn, W. (2016). Assessing Financial Education: Evidence from Boot Camp. *American Economic Journal: Economic Policy* 8(2), 322-343.
- [58] Urban, C., Schmeiser, M., Collins, J.M., & Brown, A. (2020). The Effects of High School Personal Financial Education Policies on Financial Behavior. *Economics of Education Review*, 78, 101786.
- [59] Wechsler, H., Nelson, T.F., Lee, J.E., Seibring, M., Lewis, C., & Keeling, R.P. (2003). Perception and Reality: A National Evaluation of Social Norms Marketing Interventions to Reduce College Students' Heavy Alcohol Use, *Journal of Studies on Alcohol* 64, 484-494.

A Online Appendix

A.1 Sample Selection

As mentioned in Section 2, we selected neighborhoods where households are most similar in terms of the type of houses, the age distribution, and a proxy for the income distribution. We first selected neighborhoods where at least 70% of the houses is built after 2000 using data from Statistics Netherlands (which are at the four-digit zip code level). Next, we used ING Netherlands’s client data to calculate for these neighborhoods the ratio of the 25th and 75th percentile of age and monthly inflow of money, respectively. We selected the neighborhoods that ranked among the 65% most similar for both variables. The 65th percentile was chosen as cutoff so as to get close to the target number of 40,000 emails as discussed in the main text. Finally, we dropped all neighborhoods with less than 25 households with an ING bank account, so as to guarantee anonymity. This way, we ended up with 1,904 neighborhoods with a total of 41,602 households that have a savings buffer with ING that is smaller than the median savings buffer of ING clients in the neighborhood they live in.

Since more than 15,000 clients opened the email,²⁵ we needed to make a further selection to meet the requirement of the ING Data Protection Board on the maximum number of clients we could include in our analysis. Therefore, among the clients who opened the email, we selected all those who satisfied the criteria outlined above exactly one month before the intervention as well as exactly two months before the

²⁵A total of 18,470 clients opened the email, which amounts to 44.4% of all emails that were sent.

intervention (this applies to 13,303 households). In addition we selected clients who satisfied the criteria exactly one month before the intervention, but not two months before the intervention. From this group, we selected the 1,697 clients who were the first to open the email. Our results are similar when we do not include the latter group in our analysis, see Table A2 in the next subsection.

A.2 Heterogenous Treatment Effects

Table A1: Effect of the social norm nudge 3 months after treatment

Subsample: Households that open the email in the first five days after sending it

	Buffer savings	Automatic savings amount	Automatic savings transactions
Treatment	-69.25 (86.13)	-0.87 (1.59)	-0.003 (0.004)
Dependent variable at t=0	0.42*** (0.01)	1.18*** (0.01)	0.779*** (0.005)
Constant	1608.05*** (64.85)	9.20*** (1.14)	0.035*** (0.003)
Number of households	14,362	14,362	14,362
R ²	0.10	0.56	0.62

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively.

Table A2: Effect of the social norm nudge 3 months after treatment

Subsample: Households that satisfy the selection criteria both 1 and 2 months before the treatment

	Buffer savings	Automatic savings amount	Automatic savings transactions
Treatment	-31.45 (77.14)	-0.94 (1.58)	-0.003 (0.004)
Dependent variable at t=0	1.00*** (0.03)	1.16*** (0.01)	0.778*** (0.005)
Constant	512.94*** (70.84)	9.10*** (1.14)	0.034*** (0.003)
Number of households	13,297	13,297	13,297
R ²	0.10	0.54	0.61

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively.

Table A3: Effect of the social norm nudge 3 months after treatment

Subsample: Top half of the most homogenous neighborhoods

	Buffer savings	Automatic savings amount	Automatic savings transactions
Treatment	-17.70 (123.46)	-2.29 (2.25)	-0.001 (0.005)
Dependent variable at t=0	0.27*** (0.01)	1.18*** (0.01)	0.778*** (0.007)
Constant	2015.32*** (91.92)	11.12*** (1.61)	0.035*** (0.004)
Number of households	7,979	7,979	7,979
R ²	0.05	0.56	0.62

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively.

Table A4: Effect of the social norm nudge on buffer savings 3 months after treatment

Subsamples: households distinguished by the difference between the total inflow of money and the automatic transfers in the month before the treatment.

Difference between total inflow and automatic transfers:	at least 1,000	at least 1,500	at least: 2,000	at least: 2,500
Treatment	-134.85 (96.06)	-152.80 (111.03)	-176.64 (128.87)	-143.35 (150.15)
Buffer savings at t=0	0.40*** (0.01)	0.39*** (0.01)	0.38*** (0.01)	0.37*** (0.01)
Constant	1785.44*** (72.59)	1917.05*** (83.82)	2082.89*** (97.22)	2215.32*** (113.70)
Number of households	12,112	10,223	8,538	7,110
R ²	0.10	0.10	0.09	0.09

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

Table A5: Effect of the social norm nudge on buffer savings 3 months after treatment

Subsamples: households distinguished by buffer savings 3 days before the treatment

Buffer savings 3 days before treatment:	non-negative	less than 10,000
Treatment	-71.27 (83.78)	-67.12 (82.22)
Buffer savings at t=0	0.41*** (0.01)	0.45*** (0.01)
Constant	1605.86*** (63.33)	1545.53*** (62.14)
Number of households	14,142	14,931
R ²	0.10	0.10

Standard errors are in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

Table A6: Effect of the social norm nudge on buffer savings 3 months after treatment

Subsamples: household size

Household size:	1 person	2 persons	3 or more
Treatment	-111.49 (141.41)	43.12 (147.90)	-379.78* (197.48)
Buffer savings at t=0	0.79*** (0.03)	0.41*** (0.02)	0.22*** (0.02)
Constant	755.64*** (110.57)	1668.50*** (112.00)	2525.62*** (146.37)
Number of households	5,203	5,552	4,215
R ²	0.13	0.09	0.09

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

Table A7: Effect of the social norm nudge on buffer savings for the subsample of males

	Buffer savings at t=4	Buffer savings at t=13	Buffer savings at t=38
Treatment	11.29 (91.79)	-132.92 (113.92)	-163.44 (182.29)
Buffer savings at t=0	0.91*** (0.01)	0.36*** (0.01)	0.83*** (0.02)
Constant	793.06*** (69.68)	1851.32*** (86.49)	1259.38*** (138.40)
Number of households	8,333	8,332	8,324
R ²	0.42	0.07	0.13

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

Table A8: Effect of the social norm nudge on buffer savings for the subsample of females

	Buffer savings at t=4	Buffer savings at t=13	Buffer savings at t=38
Treatment	-161.04 (87.01)	21.03 (120.11)	264.32 (193.75)
Buffer savings at t=0	0.88*** (0.01)	0.49*** (0.01)	0.72*** (0.02)
Constant	816.71*** (64.95)	1295.48*** (89.67)	1072.39*** (144.64)
Number of households	6,640	6,640	6,636
R ²	0.50	0.14	0.12

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

Table A9: Effect of the social norm nudge on buffer savings for the subsample age 48 or below

	Buffer savings at t=4	Buffer savings at t=13	Buffer savings at t=38
Treatment	55.08 (87.05)	99.60 (105.05)	115.20 (161.37)
Buffer savings at t=0	0.91*** (0.01)	0.73*** (0.02)	0.79*** (0.02)
Constant	624.39** (66.80)	774.89*** (80.63)	902.79*** (123.82)
Number of households	7,647	7,646	7,643
R ²	0.38	0.22	0.12

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

Table A10: Effect of the social norm nudge on buffer savings for the subsample age 49 or above

	Buffer savings at t=4	Buffer savings at t=13	Buffer savings at t=38
Treatment	-184.12* (93.92)	-216.14* (125.15)	-56.36 212.83)
Buffer savings at t=0	0.89*** (0.01)	0.24*** (0.01)	0.77*** (0.02)
Constant	984.00*** (69.81)	2216.67*** (93.01)	1460.78*** (158.20)
Number of households	7,341	7,341	7,332
R ²	0.50	0.04	0.13

Standard errors in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, ***, respectively.

A.3 Survey Experiment

As discussed in Section 4, we ran a small-scale survey experiment prior to the field experiment. One of the aims was to learn whether the nudge is actually noticed. For this purpose, a software developed by DVJ Insights was used that instructs the respondent to move the cursor to the position on the screen where one is looking. Figure A.1 shows a heatmap of the pictures in the treatment and control email, summarizing the attention paid by respondents to the pictures. Figure A.2 shows the percentages of respondents looking at different parts of the email using the same data. The figures are discussed in Section 4 of the paper.




Figure A.1: Heatmap summarizing how much attention participants pay to several parts of the picture in the control and treatment group

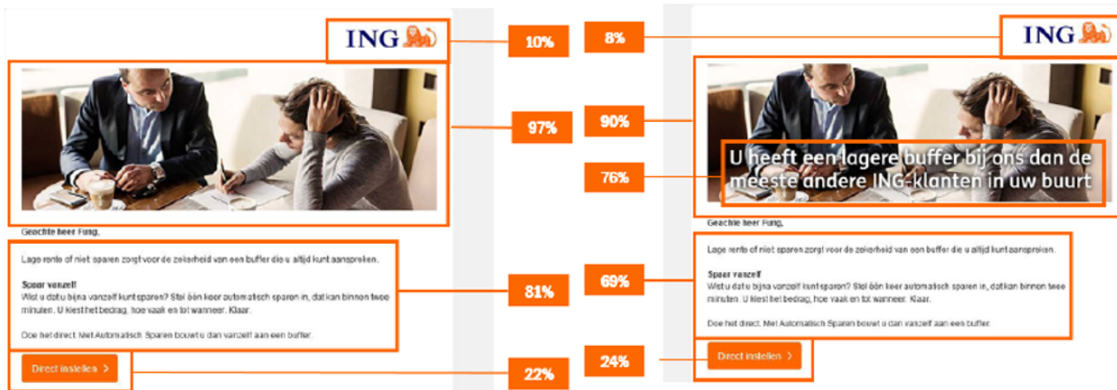


Figure A.2: Attention paid to several parts of the email in control and treatment group