

A SOCIAL NORM NUDGE TO SAVE MORE



Technical Report

Robert Dur
Dimitry Fleming
Marten van Garderen
Max van Lent

THINK
FORWARD INITIATIVE

A social norm nudge to save more: A field experiment at a retail bank *

TECHNICAL REPORT

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

September, 2019

Abstract

A large fraction of households have very little savings buffer and are therefore vulnerable to financial shocks. We examine whether a social norm nudge can stimulate such households to save more by running a small-scale survey experiment and a large-scale field experiment at a retail bank in the Netherlands. The survey experiment shows that a social norm nudge increases intended savings. In line with this, we find in our field experiment that households who are exposed to the social norm nudge click more often on a link to a personal webpage 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.

Keywords: Household Savings, Field Experiment, Nudges, Social Norms.

* This report has been prepared by the authors for the Think Forward Initiative.

† Dur: Erasmus University Rotterdam, Tinbergen Institute, CESifo, and IZA.(dur@ese.eur.nl), Fleming: ING Economics Bureau, ING Netherlands (dimitry.fleming@ing.com), van Garderen: ING Economics Bureau, ING Netherlands (marten.van.garderen@ing.com), van Lent: Leiden University (m.van.lent@law.leidenuniv.nl)

1. Introduction

A disturbingly 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 has a buffer that is too low according to the Dutch Institute for Budgetary Research and Education (Nibud 2017). Households with too little savings are at risk of having to take up expensive loans and defaulting.

Stimulating 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 expensive³. 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 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 stimulate 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 behaviour. 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, that intends to promote savings by households. We estimate the causal effect of the nudge by comparing savings of households who received the email including the nudge with savings of households who received an otherwise identical email without the nudge. 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).

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

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

⁴ See 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), Bott et al. (2019), and Giaccherini et al. (2019). However, norm-nudge interventions are not always successful, see among others Blumenthal et al. (2001), Fellner et al. (2013), Cranor et al. (2018), John and Blume (2018), and Dimant et al. (2019).

⁵ The 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.

⁶ See e.g. Duflo and Saez (2002, 2003), Hong et al. (2004), Brown et al. (2008), Kuhn et al. (2011), Brown and Laschever (2012), and Ouimet and Tate (2019).

⁷ Households in the neighborhood have also been used as a peer group in the social norm nudges that 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. Finally, note 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).

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 savings⁸.

We obtain the following results. Using a small-scale survey experiment conducted in December 2017, we first establish that the nudge attracts attention (measured using a software designed to track eye movements). Moreover, we find that the nudge increases intended savings and stimulates households to change their savings method. However, the nudge also leads to more annoyance. Next, in January 2018, ING Netherlands sent a marketing email promoting savings to more than 40,000 clients, including the social norm nudge in a random half of these. In line with the results of the survey experiment, 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 August 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 of automatic savings transactions. Our null findings are quite precisely estimated.

We run a range of robustness checks and conclude that our findings are robust. For instance, we find very

similar effects 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 Lastly, results are 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.

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 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.

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 –

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

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

¹⁰ For 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 for more details section 2.1.

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. 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 Rodriguez 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. The reference group in their intervention are co-workers of about the same age, but with sometimes very different wages. The negative effect mainly arises because the low-wage workers reduce their savings. 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. In our paper, we give social norm nudging a second chance 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.

The paper proceeds as follows. In the next section we describe the design of the field experiment. Section 3 describes the set-up and reports the results of our survey experiment. Section 4 describes the field-experimental data and section 5 our empirical strategy. Section 6 reports the results of the field experiment. Section 7 concludes.

2. Experimental Design

The field experiment took place at ING Netherlands, a large retail bank that has more than 8 million clients in the Netherlands (i.e. nearly half of the Dutch population). We set up an email marketing campaign aimed at encouraging households with little savings buffer to save more. We focus on households with a savings buffer smaller than the median savings buffer in the neighborhood they live in. Moreover, as described below in more detail, we exclude households that seem to have little opportunity to increase their savings.

ING Data Protection Board put a maximum of 15,000 on the number of clients for which we could analyze detailed anonymized data, implying that we needed to make a choice whom to target. Inspired by the results in Beshears et al. (2015), we decided to focus on households who live in relatively homogeneous neighborhoods. This way, we aim to minimize the discouragement effect arising from upward social

comparison that was found in Beshears et al. (2015). We describe the sample selection in detail in subsection 2.1.

Households in the sample were randomly assigned to treatment and control. Following the approach in Cr'epon et al. (2013), we randomly varied the fraction of households assigned to treatment across regions, so as to be able to detect spillover effects. The exact randomization procedure is described in subsection 2.2.

Households in the treatment group receive an email containing the social norm nudge, those in control receive an otherwise identical email without the nudge.¹¹ The social norm nudge is displayed in the picture in the treatment email and reads: "You have a lower buffer with us than most other ING clients in your neighborhood". Figure 1 shows the control email (left-hand side) and treatment email (right-hand side).

Figure 1: The email sent to households in the control group (left picture) and in the treatment group (right picture)



¹¹ Note that we chose not to use a control group of clients who receive no email at all. While it would be interesting to also be able to study the effect of receiving the email as compared to not receiving an email, the maximum on the number of clients we could study imposed by ING Data

Protection Board meant we needed to make choices. A power analysis revealed we might lack power to precisely estimate treatment effects had we decided to split the group in more than two subgroups. Hence we decided to limit our study to the test of the social norm nudge.

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

"Dear Mr. [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").

2.1 Sample selection

As we mentioned above, ING Data Protection Board allowed us to analyze detailed anonymized microdata on a maximum of 15,000 clients. Since the company knows from previous promotional campaigns that about 40% of the clients open marketing emails from ING, we targeted roughly 40,000 clients.

The social norm nudge in the treatment email makes a comparison of the client's savings with savings of others in the same neighborhood. We define neighborhoods as five-digit zip code areas, implying that each neighborhood has on average about

250 households (of which about half are ING clients). The Netherlands consists of 33,000 of such neighborhoods.

To avoid the discouragement effect found in Beshears et al. (2015), 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. 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 containing 343,088 households.

We subsequently selected, for these neighborhoods, those households who have below median buffer savings, where buffer savings are defined as the sum of the amount on the current account and amounts on the liquid savings accounts. We excluded households with a negative buffer. Moreover, we excluded households who were most likely not able to increase their savings. That is, we imposed that a household should have a sufficient regular inflow of money of at least 1,000 euros per month. This requirement also makes it likely that we select clients for whom their ING account is their primary bank account.¹² Lastly, we also 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. This leaves us with a sample of 41,602

¹² Clearly, 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 (DNB Household Survey), a majority of ING clients do not hold a current account at another bank.

households, who were all sent either a control or treatment email.

The sample we analyze consist of the households that opened the email and loaded the picture that is in the email, which is tracked by ING Netherlands. Since the treatment message was included in the picture, and all other parts of the email (including the subject line) are identical between treatment and control, there can be no selection into treatment.¹³

Slightly more than 15,000 clients opened the email, so 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 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.¹⁴ The clients we selected live in 1,904 different neighborhoods.

2.2 Randomization and Spillovers

Our experimental design allows us to detect possible contamination of the control group following the approach set out in Crépon et al. (2013). As discussed in the Introduction, spillover effects may occur when control households learn about saving plans or observe consumption patterns of treated households. We expect potential spillovers to be the largest within the neighborhood. Therefore, we varied treatment intensity by neighborhood. In a random half of the neighborhoods, we assigned 80% to treatment and 20% to control. In the other half, we 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 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.

¹³The subject line in both treatment and control was: "What if things ever go wrong?".

¹⁴ In subsection 6.2 we show that none of our results change when we drop this part of the sample.

3. Survey Experiment

The large-scale field experiment was preceded by a small-scale survey experiment in order to study the following three issues: i) Does the nudge attract the clients' attention?; ii) Do clients increase their intended buffer savings in response to the nudge?; and iii) Is the nudge perceived as intrusive by clients?¹⁵ 292 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 field experiment: clients at ING, with low savings, sufficient income, and living in homogeneous neighborhoods. Each survey participant was randomly selected to see one of the messages and was subsequently asked questions about his or her motivation to save, the intrusiveness of the message, as well as his or her perceptions about ING.

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 2 shows a heatmap of the pictures in the treatment and control email, summarizing the attention paid by respondents to the pictures. Clearly, respondents in the treatment group spend more time looking at the picture, and particularly so at the part that displays the social norm nudge. Figure 3 shows the percentages of respondents looking at different parts of the email using the same data. 76% of the respondents in the treatment group reads the social norm nudge. Figure 3 also makes clear that 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 makes clear to some respondents that the email is not relevant to them. 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 stimulated 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.

Figure 2: Heatmap summarizing how much attention participants pay to several parts of the picture in the control and treatment group.



¹⁵ The survey experiment included, in addition to our treatment email and control email, two other emails with

other versions of the social norm, which we ignore here for brevity.

Figure 3: Attention paid to several parts of the email in control and treatment group

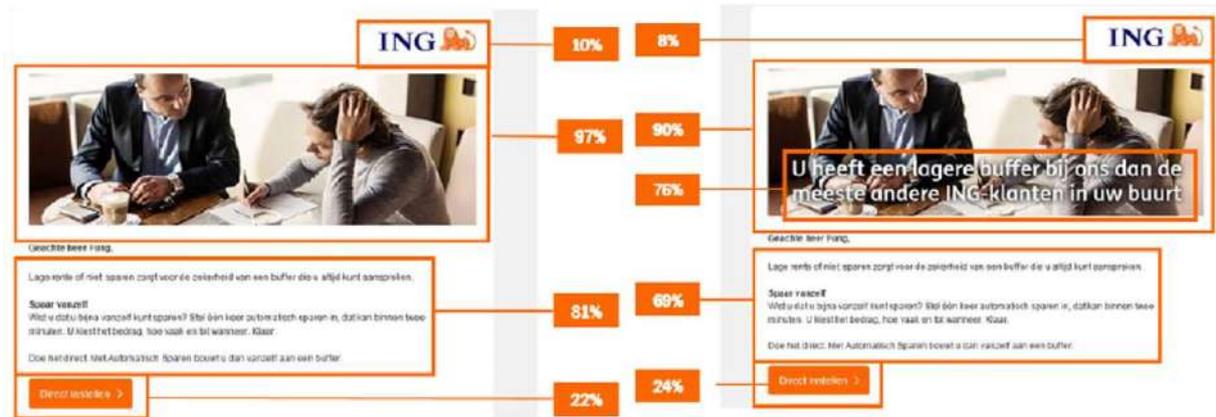


Table 1 provides support for this interpretation. It shows the mean of answers given to a range of questions of respondents in the treatment and control group.

Clearly, respondents in the treatment group find the email more relevant (22% versus 13%). It is also clear from the table that 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 1 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%).¹⁶ Despite this, there do not seem to be major repercussions for the image of the bank, see the bottom part of Table 1.¹⁷

¹⁶ 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).

¹⁷ Note 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 1 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 agreeing with the statement (i.e. score 6 or higher on a 7-point-scale).

4. Field Experiment: Descriptives

We use anonymized micro-level data on 15,000 households. Our key variable of interest is the households' buffer savings, which we define as the total euro amount on its current accounts and savings accounts (those without withdrawal limits). We have weekly data (the amounts on Sunday each week). We also know the number of automatic savings transactions in each week. Further, we obtained some demographics, such as age and household size. We also know how much money flows into and out of the bank accounts. Finally, we have data on whether clients clicked on the links in the email.

Table 2 provides the descriptive statistics measured on Sunday 31 December 2017, a few days before the treatment. Households in our sample have on average a buffer of about 2,000 euros (the median is about 1,400 euros).¹⁸ During the week before treatment, about one out of six households made an automatic

savings transaction and the average amount of automatic savings across all households is about 23 euros. None of these savings behaviors differ significantly between the treatment and control group. The same holds for our demographic variables household size and age. On average a household has 2.1 members, and the average age is close to 47 years.¹⁹

Table 3 shows the immediate impact the inclusion of the social norm nudge had, as measured by clicks on the links in the email and automatic savings transactions. We find that households in the treatment group click significantly more often (3.4% versus 2.7%) on the link to their automatic savings page. This is in line – at least qualitatively – with the findings from the survey experiment in Table 1, showing increased intended savings. Table 3 also shows that clients in the treatment group are not more likely to opt out of future marketing email messages from ING (0.5% opts out of

Table 2: Descriptive statistics the last Sunday before treatment (12/31/2017)

	Control N=7,458	Treatment N=7,542	p-value
Current + savings account	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
Age	46.81 (14.33)	46.48 (14.58)	0.13

Standard deviations are in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The current + savings account is the total amounts in all current and savings accounts. The automatic savings amount is the total amount transferred to a savings account in the last week. Current + savings account and automatic savings amount are measured in euros. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction during the last week.

¹⁸ Note that the standard deviations are substantial, which is caused by a limited number of outliers.

¹⁹ To be precise, we only observe household members and their age if they have an account at ING.

Table 3: Descriptive statistics the first Sunday after treatment (01/07/2018)

	Control N=7,458	Treatment N=7,542	p-value
Click on automatic savings page	0.027 (0.162)	0.034 (0.180)	0.00***
Opted out of future emails	0.005 (0.069)	0.005 (0.069)	0.994
Automatic savings transaction	0.090 (0.287)	0.086 (0.281)	0.38

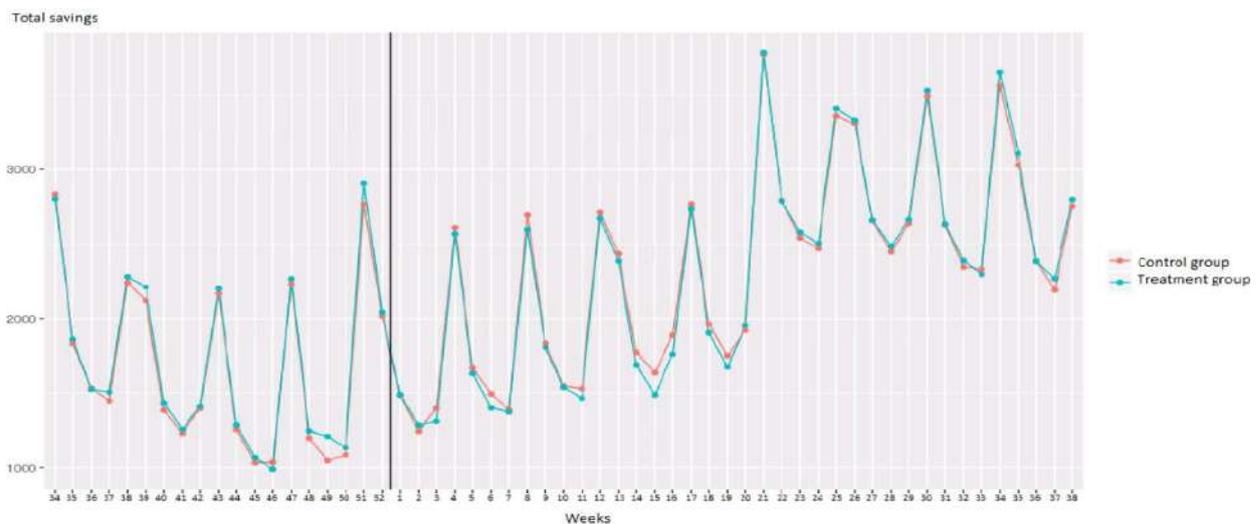
Standard deviations are in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The click on automatic savings page and the opted out of future emails are dummy variables. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction during the last week.

ING's email-list in both control and treatment group). This contrasts the expectations raised by the survey experiment, which showed significant increases in annoyance in response to the email including the social norm nudge, see Table 1.²⁰ The fraction of households that made at least one automatic savings transaction in the week of the treatment does not differ between treatment and control group.²¹ Despite this lack of an immediate effect on automatic savings transactions, there may be an effect on automatic savings later on, as households often sign up for an automatic savings

plan that does not start immediately but at a future point in time.

In Figures 4, 5, and 6, we take a first look at the treatment effects in the long run. The figures show the average buffer of households (Figure 4), the percentage of households making an automatic savings transaction (Figure 5), and the average amount saved automatically (Figure 6) for the treatment and control group separately for all weeks before and after the treatment.

Figure 4: 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.



²⁰ In a field experiment with a charity, Damgaard and Gravert (2018) find that nudging increases unsubscriptions from the mailing list.

²¹ Note that the shares are much lower in both treatment and control as compared to the week before, as shown in Table 2. This reflects a within-monthly pattern that is visible each month, see Figure 5 below

Figure 5: Share of households making an automatic savings transaction in the treatment and control group. Note: the black vertical line denotes the time the email was sent.

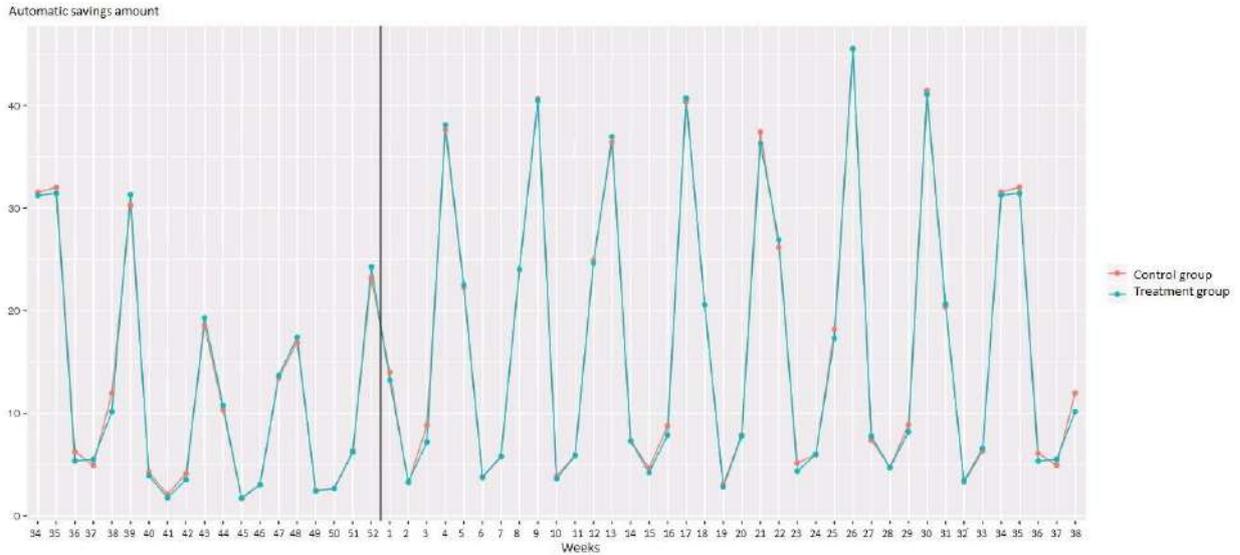
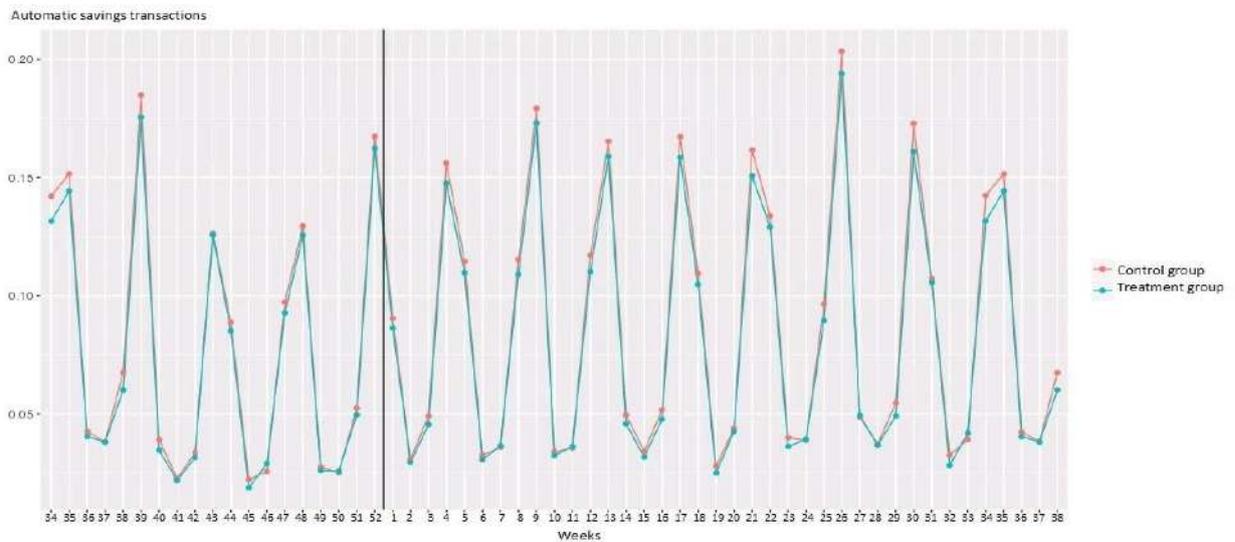


Figure 6: Average amount of automatic savings transactions in euros in the treatment and control group. Note: the black vertical line denotes the time the email was sent.



The black vertical line indicates the timing of the treatment. It is clear that the savings variables vary 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. The peaks in Figure 4 at the end of the year (two weeks before treatment) and in May (the fifth spike after treatment) are due to holiday pay, which is commonly paid out in December and May in the Netherlands. Further, we see before as well as after

the treatment only minor differences in the averages for control and treatment group households, indicating that the randomization was successful and that there is no average treatment effect of the social norm nudge. The same is true for the frequency and size of automatic savings transactions, see Figures 5 and 6.

5. Field Experiment: Empirical Strategy

5.1 Main Specifications

Our data run from the end of August 2017 to the end of September 2018, and our treatment takes place in the first week of January 2018. We estimate an OLS panel regression containing household and week fixed effects. Our primary outcome variable is the total buffer savings amount transformed using an inverse hyperbolic sine (IHS) transformation. This transformation has the properties of the natural logarithm and marginal effects can be interpreted in the same way as the natural logarithm. An advantage of this transformation is that it allows for retaining zero-valued and negative-valued observations (see Johnson 1949 and Burbidge 1988; for recent applications using wealth data see e.g. Carroll et al. 2003 and Shaefer et al. 2013). We use the same transformation for our secondary outcome variable, the amount of automatic savings. Our third outcome variable is an indicator whether a household made at least one automatic transaction to a savings account. We thus estimate:

$$y_{it} = \alpha_i + \tau_t + \beta P_t T_i + \varepsilon_{it}, \quad (1)$$

where y_{it} is the outcome variable, i is a household, t denotes the week, P_t equals zero before and one in the weeks after the intervention, T_i is a dummy that equals one for treatment households and zero for households in the control group. The coefficient of interest is β , the estimate of the effect of the social norm nudge. Finally, ε_{it} is the residual. Throughout, we cluster standard errors at the household level.

5.2 Dynamic Effects

We estimate the dynamic effects of the social norm nudge by interacting the treatment dummy T_i with time dummies. Specifically, we estimate:

$$y_{it} = \alpha_i + \tau_t + \sum_{t=1}^Z \beta_t T_i + \varepsilon_{it}, \quad (2)$$

where T_i is a dummy that equals one if household i belongs to the treatment group and zero if household i belongs to the control group. β_t is the estimated difference between treatment and control in period t . Z includes all periods, except for the period right before the treatment, which we take as reference period. Finally, ε_{it} is the residual.

5.3 Spillover Effects

In order to examine whether spillover effects are important in our context we estimate:

$$y_{it} = \alpha_i + \tau_t + \beta_{TH} P_t T_{iH} + \beta_{TL} P_t T_{iL} + \beta_{CH} P_t C_{iH} + \varepsilon_{it}, \quad (3)$$

where T_{iH} is a dummy for households in the treatment group living in a high treatment intensity area (i.e., where 80% is assigned to treatment), T_{iL} is a dummy for households in the treatment group living in a low treatment intensity area (i.e., where 20% is assigned to treatment), and C_{iH} is a dummy for households in the control group living in a high treatment intensity area. The reference group is households in the control group living in a low treatment intensity area. The coefficient β_{CH} is expected to be zero if there are no spillover effects, implying that the households in the control group do not behave differently in the high as compared to the low treatment intensity areas.

6. Field Experiment: Results

6.1 Average Treatment Effect

Table 4 shows the average treatment effects, estimated using equation (1).²² In line with Figures 4 to 6, we find that the treatment did neither affect households' buffer savings nor whether they saved automatically. Also, the amount saved automatically is not affected by the treatment. The point estimates are very small and quite precisely estimated.

6.2 Heterogeneous Treatment Effects

In Tables A1-A5 of the Appendix, we examine whether the treatment effect differs for several subsamples, and find that our null results are robust. We study 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 differ in their interest in emails from ING and may pay more attention to it; ii) Those for whom our selection criteria apply more broadly, namely up to and including two months before the treatment (89% of the sample); iii) Those living in the top half most homogeneous neighborhoods (53% of the sample), because the social norm nudge may be considered as more relevant by subjects whose neighbors are more similar to them; 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 in the month before the intervention; and v) Those with a non-negative buffer three days before the email arrives (97% of the sample), and those whose buffer is not exceeding 10,000 euros three days before the email arrives (99.9% of the sample). 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 opportunities to save more. Our analysis of heterogeneous treatment effects exploits the remaining variation in the selected sample, which is limited.

6.3 Dynamic Effects

It may be that, although households on average do not respond over the full post-treatment period, there are interesting dynamics in their response. For instance, households may respond immediately to the nudge, but forget about the nudge shortly after. It may also be

Table 4: Effect of the social norm nudge on saving behavior

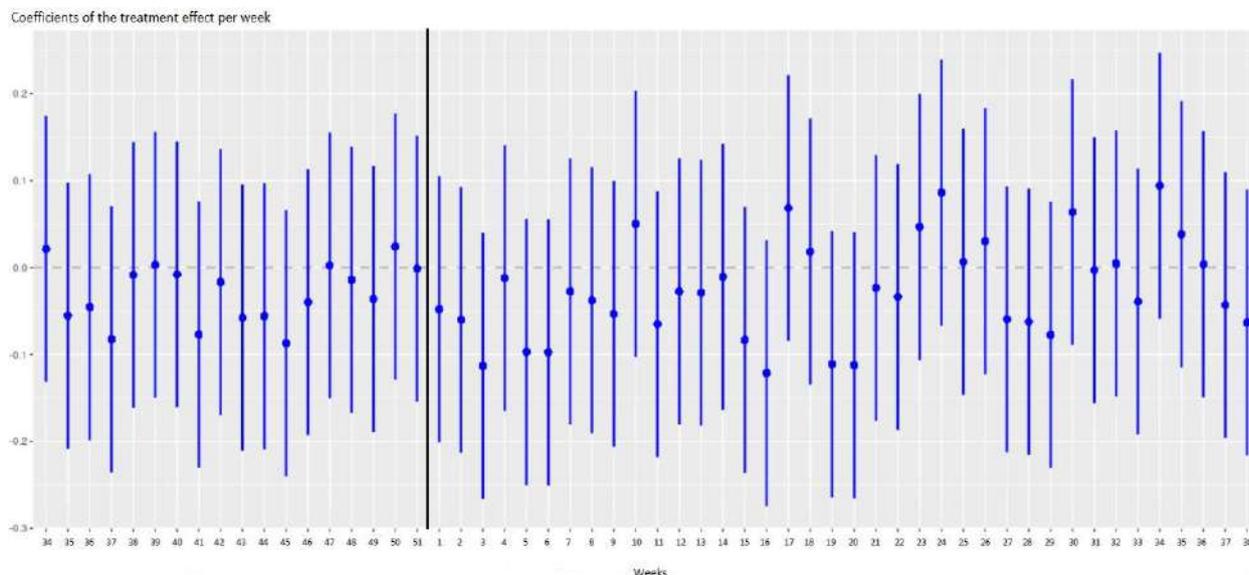
	Total savings	Automatic savings amount	Automatic savings transactions
Treatment	0.002 (0.015)	-0.006 (0.006)	-0.001 (0.001)
Week fixed effects (57 weeks)	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Number of households	15,000	15,000	15,000
Number of observations	854,793	854,793	854,793
R ²	0.088	0.056	0.053

Standard errors are clustered at the household level and in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The variables total savings and automatic savings amount are transformed using the inverse hyperbolic sine transformation. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction per week.

²² Note that, out of 855,000 potential client-week observations, we have 207 missing client-week observations. These represent clients who, at some point after treatment, no longer show up in the data because their account has been

closed. The first missing observation shows up ten weeks after the treatment. There is no effect of the treatment on attrition.

Figure 7: Dynamics in the treatment effect



Note: the black vertical line denotes the time of the treatment. Week 52 is the reference week and is therefore omitted. The vertical axis shows the estimated treatment effect on current savings accounts for each week and the confidence interval.

that it takes some time for households to adjust their spending patterns, implying that there is little or no response initially and more later on. We estimate the dynamic effect of the nudge by estimating equation (2), where a period is a week and we take the last week before treatment as the reference category. Figure 7 shows the estimated coefficients β_t and the confidence intervals of the treatment effect for total savings (the table with the regression results is available upon request). In line with Figure 4, we find no meaningful differences between treatment and control, neither before nor after the intervention.

6.4 Spillover Effect

As argued in the Introduction, contamination may result in a downward bias in the estimated treatment effect. By virtue of our design, we can shed some light on

whether control group subjects have been affected by the treatment. Following the approach in Crépon et al. (2013) we randomly varied the fraction of households who receive treatment by neighborhood. If contamination is important, we expect that savings behavior of control households depends on the fraction of treated households in their neighborhood. We find no indication for this, see the results of estimating equation (3) in Table 5. Unexpectedly, we do find a difference in response of treated households to treatment between high-dose and low-dose neighborhoods for one of our outcome variables (see the first column), but not for the other two (see the second and third column).

Table 5: Spillover effects of the social norm nudge

	Total savings	Automatic savings amount	Automatic savings transactions
Control × High-dose treatment	0.017 (0.027)	0.0002 (0.0010)	-0.0002 (0.0020)
Treatment × Low dose treatment	-0.086** (0.027)	0.003 (0.001)	0.002 (0.002)
Treatment × High dose treatment	0.028 (0.017)	-0.008 (0.006)	0.001 (0.001)
Week fixed effects (57 weeks)	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Number of households	15,000	15,000	15,000
Number of observations	854,793	854,793	854,793
R ²	0.088	0.056	0.053

Standard errors are clustered at the household level and in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The variables total savings and automatic savings amount are transformed using the inverse hyperbolic sine transformation. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction per week. The variable control × high dose treatment is a dummy that equals 1 if the household is in control in a neighborhood where 80% was in the treatment. Likewise, the variable treatment × low dose treatment is a dummy that equals 1 if a household is in treatment in neighborhood where 20% was in treatment.

7. Conclusion

We ran a small-scale survey experiment and 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 have found that adding the nudge increases households' intended savings as measured by the responses of subjects in the survey experiment. Moreover, using data on clicks and website visits from the field experiment, we have found that adding the nudge stimulates 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. The field experiment by Kast et al. (2018) that studied a similar nudge provided to microcredit clients in Chile further increased our expectations that the nudge would be effective, even

though that study could not rule out that the behavioral response was mainly due to a reminder effect. We also attempted to minimize the discouragement effect that social norm nudges can have in heterogeneous groups (Beshears et al. 2015). Yet, no substantive effect of the nudge resulted, and this null finding is quite precisely estimated. 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.

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.

8. Acknowledgments

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 also thank 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.

9. References

- Akerlof, G. A., & Kranton, R. E. (2000). Economics and Identity. *Quarterly Journal of Economics*, 115(3), 715-753.
- Allcott, H. (2011) Social Norms and Energy Conservation. *Journal of Public Economics*, 95(9-10), 1082–1095.
- 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.
- 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.
- Bernheim, B. D. (1994). A Theory of Conformity. *Journal of Political Economy*, 102(5), 841-877.
- 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.
- 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.
- 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.
- 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.
- 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.
- Bicchieri, C., Dimant, E. (2019). Nudging with Care: The Risks and Benefits of Social Information. *Public Choice*, forthcoming.
- 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.
- Bordalo, P., Gennaioli, N., Shleifer, A. (2013). Saliency and Consumer Choice. *Journal of Political Economy*, 121(5), 803-843.
- Bott, K. M., Cappelen, A. W., Sorensen, E., & Tungodden, B. (2019). You've Got Mail: A Randomised Field Experiment on Tax Evasion. *Management Science*, forthcoming.
- Bradler, C., Dur, R., Neckermann, S., & Non, A. (2016). Employee Recognition and

- Performance: A Field Experiment. *Management Science*, 62(11), 3085-3099.
- 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.
- 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.
- 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.
- 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.
- Burbidge, J. B., Magee, L., & Robb, A. L. (1988). Alternative Transformations to Handle Extreme Values of the Dependent Variable. *Journal of the American Statistical Association*, 83(401), 123-127.
- Carroll, C. D., Dynan, K. E., & Krane, S. D. (2003). Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets. *Review of Economics and Statistics*, 85(3), 586-604.
- 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.
- 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.
- 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.
- 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.
- Cranor, T., Goldin, J., Homonoff, T., Moore, & L. (2018). Communicating Tax Penalties to Delinquent Taxpayers: Evidence from a Field Experiment. Mimeo, Stanford Law School.
- 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.
- Damgaard, M.T., & Gravert, C. (2018). The Hidden Costs of Nudging: Experimental Evidence from Reminders in Fundraising. *Journal of Public Economics*, 157, 15-26.
- Dimant, E., Van Kleef, G.A., & Shalvi, S. (2019). Requiem for a Nudge: Framing Effects in Nudging Honesty. Mimeo, University of Pennsylvania.

- 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.
- 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.
- 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.
- 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.
- Giaccherini, M., Herberich, D.H., Jimenez-Gomez, D., List, J.A., Ponti, G., & Price, M.K. (2019). The Behaviorist Goes Door-To-Door: Understanding Household Technological Diffusion Using a Theory-Driven Natural Field Experiment. NBER Working Paper No. 26173.
- 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.
- 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.
- 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.
- Hallsworth, M., List, J.A., Metcalfe, R.D., & Vlaev, I. (2017). The Behaviorist as Tax Collector: Using Natural Field Experiments to Enhance Tax Compliance. *Journal of Public Economics*, 148, 14-31.
- Hong, H., Kubik, J.D., & Stein, J.C. (2004). Social Interaction and Stock-Market Participation. *Journal of Finance*, 59(1), 137-163.
- 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.
- Johnson, N.L. (1949). Systems of Frequency Curves Generated by Methods of Translation. *Biometrika*, 36(1/2), 149-176.
- 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.
- 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.
- Kast, F., Meier, S., & Pomeranz, D. (2018). Saving More in Groups: Field Experimental Evidence from Chile. *Journal of Development Economics*, 133, 275-294.

- 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.
- 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-geldhand/>
- Ouimet, P., & Tate, G. (2019). Learning from Coworkers: Peer Effects on Individual Investment Decisions. *Journal of Finance*, forthcoming.
- 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.
- 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.
- Skimmyhorn, W. (2016). Assessing Financial Education: Evidence from Boot Camp. *American Economic Journal: Economic Policy* 8(2), 322-343.
- Shaefer, H. L., Song, X., & Shanks, T.R.W. (2013). Do Single Mothers in the United States Use the Earned Income Tax Credit to Reduce Unsecured Debt? *Review of Economics of the Household* 11(4), 659-680.
- Urban, C., Schmeiser, M., Collins, J.M., & Brown, A. (2018). The Effects of High School Personal Financial Education Policies on Financial Behavior. *Economics of Education Review*, forthcoming.
- 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.

10. Appendix

Table A1: Effect of the social norm nudge on saving behavior

Sample: Group of households who open the email in the first five days after sending it

	Total savings	Automatic savings amount	Automatic savings transactions
Treatment	0.006 (0.015)	-0.005 (0.006)	-0.0001 (0.001)
Week fixed effects (57 weeks)	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Number of households	14,363	14,363	14,363
Number of observations	818,484	818,484	818,484
R ²	0.089	0.056	0.054

Standard errors are clustered at the household level and in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The variables total savings and automatic savings amount are transformed using the inverse hyperbolic sine transformation. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction per week.

Table A2: Effect of the social norm nudge on saving behavior

Sample: Group of households who satisfied the selection criteria in both 1 and 2 months before the experiment

	Total savings	Automatic savings amount	Automatic savings transactions
Treatment	-0.004 (0.015)	-0.006 (0.006)	-0.0001 (0.001)
Week fixed effects (57 weeks)	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Number of households	13,303	13,303	13,303
Number of observations	758,094	758,094	758,094
R ²	0.093	0.055	0.053

Standard errors are clustered at the household level and in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The variables total savings and automatic savings amount are transformed using the inverse hyperbolic sine transformation. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction per week.

Table A3: Effect of the social norm nudge on saving behavior

Sample: Top half of the most homogenous neighborhoods

	Total savings	Automatic savings amount	Automatic savings transactions
Treatment	0.005 (0.020)	0.002 (0.008)	0.0001 (0.001)
Week fixed effects (57 weeks)	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Number of households	7,979	7,979	7,979
Number of observations	454,641	454,641	454,641
R ²	0.090	0.061	0.059

Standard errors are clustered at the household level and in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The variables total savings and automatic savings amount are transformed using the inverse hyperbolic sine transformation. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction per week.

Table A4: Effect of the social norm nudge on saving behavior

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

	full sample	min: 1,000	min: 1,500	min: 2,000	min: 2,500
Dependent variable: Total savings					
Treatment	0.002 (0.015)	0.012 (0.016)	0.015 (0.018)	0.004 (0.020)	-0.016 (0.023)
Week fixed effects (57 weeks)	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes
Number of households	15,000	12,367	10,189	8,132	6,454
Number of observations	854,793	704,797	580,684	463,504	367,87
R ²	0.088	0,088	0,086	0.084	0.082

Standard errors are clustered at the household level and in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The variables total savings and automatic savings amount are transformed using the inverse hyperbolic sine transformation. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction per week. In columns 2 to 5 we split up the sample, and focus on those households who have a difference between the total inflow of money and the automatic transfers of at least 1,000, 1,500, 2,000, and 2,500 respectively.

Table A5: Effect of the social norm nudge on saving behavior

Sample: households distinguished by the total savings 3 days before the experiment

Sample: total savings 3d before the experiment	full sample	non-negative	less than 10,000
Dependent variable: Total savings			
Treatment	0.002 (0.015)	0.010 (0.014)	0.010 (0.015)
Week fixed effects (57 weeks)	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Number of households	15,000	14,578	14,983
Number of observations	854,793	830,744	853,824
R ²	0.088	0.089	0.088

Standard errors are clustered at the household level and in parentheses, p-values < 0.10, 0.05, 0.01 represented by *, **, *** respectively. The variables total savings and automatic savings amount are transformed using the inverse hyperbolic sine transformation. The automatic savings transaction variable is a dummy variable that equals 1 in case of at least one automatic savings transaction per week.

The authors



Robert Dur
Department of Economics, Erasmus
University Rotterdam, Tinbergen
Institute, CESifo, and IZA
E-mail: dur@ese.eur.nl



Dimitry Fleming
ING Economics Bureau, ING
Netherlands
E-mail: dimitry.fleming@ing.com



Marten van Garderen
ING Economics Bureau, ING
Netherlands
E-mail: marten.van.garderen@ing.com



Max van Lent
Department of Economics, Leiden
University
E-mail: m.van.lent@law.leidenuniv.nl

Disclaimer

This 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 Program is to inspire practical research insights in the financial decision-making domain. It does not constitute any financial advice or service offer.