



UNIVERSITY OF GOTHENBURG
SCHOOL OF BUSINESS, ECONOMICS AND LAW

WORKING PAPERS IN ECONOMICS

No 650

**The hidden costs of nudging: Experimental evidence
from reminders in fundraising**

Mette Trier Damgaard & Christina Gravert

March 2016

ISSN 1403-2473 (print)
ISSN 1403-2465 (online)

The hidden costs of nudging: Experimental evidence from reminders in fundraising*

Mette Trier Damgaard[†]
Aarhus University

Christina Gravert[‡]
University of Gothenburg

March 14, 2016

Abstract

We document the hidden costs of one of the most policy-relevant nudges, reminders. Sending reminders, while proven effective in facilitating behavior change, may come at a cost for both senders and receivers. Using a large scale field experiment with a charity, we find that reminders increase donations, but they also substantially increase unsubscriptions from the mailing list. To understand this novel finding, we develop a dynamic model of donation and unsubscription behavior with limited attention which is tested in reduced-form using a second field experiment. We also estimate our model structurally to perform a welfare analysis, showing that reminders are welfare diminishing for the potential donors as non-givers incur a welfare loss of \$2.35 for every reminder. The net benefit of every reminder to the charity is \$0.18. Our evaluation shows the need to evaluate nudges on their intended as well as unintended consequences.

Keywords: Avoiding-the-ask, charitable giving, field experiment, inattention, nudge, reminders.

JEL codes: C93, D03, D64, H41

*We are grateful to Stefano DellaVigna for encouraging this research and for numerous discussions and comments. This research would not have been possible without the generous cooperation of DanChurchAid and especially Kim Haakansson, Nina Halberstadt, Allan Lindemark, and Ole Dahl Rasmussen. We also thank Martin M. Andreasen, Ned Augenblick, Teodora Boneva, Alexander Koch, John List, Ulrike Malmendier, Takeshi Murooka, Georg Nöldeke, Heiner Schuhmacher, Robert Sudgen, seminar participants in Berkeley, Gothenburg, and Aarhus, and participants at the 2015 European ESA meeting, 2015 San Diego Spring School in Behavioral Economics, 2015 Spring CNEE Workshop, 2015 DGPE Workshop, 2015 ESA Mentoring Workshop, 2015 LPEX Workshop for comments and suggestions. This research was partly funded by the Swedish Research Foundation and Aarhus University. M. Damgaard thanks Købmand Ferdinand Sallings Mindefond for financial support and UC Berkeley for hospitality in the academic year 2014/15. This RCT was registered in the American Economic Association Registry for RCTs under trial number 762.

[†]Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark, e-mail: mdamgaard@econ.au.dk.

[‡]Department of Economics, University of Gothenburg, Room D606, Vasagatan 1, SE 41124 Gothenburg, Sweden, e-mail: christina.gravert@economics.gu.se

1 Introduction

In recent years, “nudging” policies have gained increased attention both from practitioners and from academics. Nudges are small deliberate changes to the decision environment designed to increase privately and socially beneficial behavior such as healthy habits, savings, environmental protection, or charitable giving without altering prices or taking away any options. Nudging interventions often require low implementation costs and induce significant positive behavioral change, which has fueled enthusiasm for its use among policy makers. However, evaluating the success of a nudge on the magnitude of behavioral change and implementation cost alone could be misleading from a social welfare perspective.

This paper provides a theoretical and empirical analysis of one of the most frequently applied and well-known nudges: reminders. Reminders are designed to bring a particular decision or task to recipients’ attention and as a result induce behavioral change. A large number of recent papers have shown that reminders can influence behavior in the context of gym attendance (Calzolari and Nardotto, 2014), adherence to medical treatments (Vervloet et al., 2012; Altmann and Traxler, 2014), personal savings (Karlán et al., 2012), take-up of social benefits (Bhargava and Manoli, 2015), electricity consumption (Allcott and Rogers, 2014; Gilbert and Zivin, 2014), and giving to charitable organizations (Huck and Rasul, 2010; Sonntag and Zizzo, 2015).

Technological improvements over the past few decades have led to low distributional cost of reminders, implying that reminders are likely to become even more common in coming years. This makes it relevant to explore whether there are costs to using reminders.

Reminders impose time and effort costs on recipients as well as psychological costs such as annoyance and guilt. These costs should be taken into account when evaluating the nudge and could constitute a “hidden cost of nudging”. A cost inducing reminder only *increases* the utility of the recipient if it prompts a behavioral change. Hence, if many people are nudged, but only few change their behavior, then the welfare effects of the nudge could be negative. The annoyance costs could lead recipients to opt out of the reminder and block the communication channel.

In addition to the potential and immediate costs for recipients, reminders can also lead to long-term costs for the sender if recipients unsubscribe from the mailing list. The sender therefore has to weigh current returns against long-term costs due to the loss of subscribers. Depending on his discount rate, the costs might outweigh the benefits despite the low distributional costs.

We examine the reminder effect, including the hidden costs, in the context of charitable giving.

In particular, we consider a revealed preference measure of the “annoyance costs”: unsubscribing from reminder messages.

To simultaneously understand giving and unsubscription behavior, we develop a dynamic model of warm-glow giving where individuals incur an annoyance cost every time the charity sends a fundraising appeal. Annoyance costs can be psychological costs such as guilt or perceived pressure or practical costs such as time and attention. Every period, individuals decide whether to give or not if reminded about the donation possibility (Andreoni, 1989, 1990; DellaVigna et al., 2012). In addition, individuals have the option to unsubscribe from future communication, making dynamic considerations relevant. By unsubscribing they avoid future annoyance costs associated with reminder messages, but they also risk missing future opportunities to donate. We model forgetting by incorporating inattention similar to Karlan et al. (2012); Taubinsky (2013), and Calzolari and Nardotto (2014). Our model predicts a higher rate of giving and a higher unsubscription rate in response to reminders. We show that the unsubscription decision further depends on whether people evaluate the option value of staying subscribed to be sufficiently large to justify anticipated future annoyance costs.

We test these predictions in two field experiments with a charity. In field experiments, participants are not aware that their behavior is being observed, and we therefore observe their natural reactions. The first experiment tests the prediction that reminders increase unsubscriptions by sending solicitation e-mails to approximately 17,000 warm-list donors, i.e., individuals who have donated to the charity in the recent past. Individuals in the control group receive one e-mail asking them to donate within ten days. People in the treatment group receive the same e-mail and an additional reminder one week later. In line with the predictions of the model, we find that the reminder significantly increases donations but also significantly increases unsubscriptions from the mailing list.

The second field experiment tests the prediction that the unsubscription choice is determined by the option value of subscribing and anticipated annoyance costs. A sample of 43,000 previous donors receives a regular solicitation e-mail from the charity. With the exception of one sentence, the solicitation e-mails are identical across our three treatment groups. In the control treatment, potential donors are made aware that the charity sends an e-mail to individuals on the mailing list approximately once per month. This is to fix expectations of how often people should expect to receive e-mails from the charity, and it allows potential donors to form beliefs about future annoyance costs. In our Low Frequency treatment, we exogenously decrease anticipated annoyance

costs relative to the control treatment by announcing that the charity will only send *one* e-mail in the next three months. The model predicts that individuals in the Low Frequency group are less likely to unsubscribe than individuals in the control group. In the Future Benefit treatment, we increase the option value of staying on the list by announcing that next month an anonymous donor will make a donation for every person on the mailing list who donates in response to the next e-mail. Compared to the control group, this message should make individuals less willing to unsubscribe because the utility from donating in the next period is increased. Four weeks later, all participants received the same e-mail about a donation opportunity, regardless of which treatment they were assigned to.

In line with our model predictions, we find that announcing a reduced frequency of mailings significantly decreases the number of unsubscriptions relative to the control treatment. Announcing a future matching opportunity also reduces the number of unsubscriptions, but this result is only marginally significant. While the main outcome of interest is the decision to unsubscribe, we also measure how these nudges affect the decision to donate. The treatments have no effect on the decision to donate or the donated amount which is also consistent with our model.

Using the results from the second experiment, we structurally estimate the annoyance costs of being solicited through e-mail. We then use these estimates to conduct a welfare analysis from the perspective of the recipients. For the potential donor, there is a trade-off between annoyance costs and warm-glow from remembering to give. We estimate the costs associated with *receiving* a reminder to 12.95 DKK (\$2.35). This can be interpreted as the willingness to pay not to receive a reminder, such as buying a sophisticated e-mail ad-blocker. Since donations are relatively rare in our sample, this implies that a reminder on average is welfare diminishing for potential donors.

We then consider the perspective of the charity by estimating the impact of a reminder on donations, using individual level donation data from the first experiment. For the charity the trade-off is between an immediate increase in donations and a future loss of donations from unsubscribing individuals. Since we only solicit individuals who have donated to the charity in the past, the charity loses former donors, albeit donors who on average gave at slightly lower levels than the people who stay subscribed. After unsubscribing, some individuals continue to donate to the charity through other channels (i.e. text messages, bank transfers, in store donations), but at significantly lower levels than while they were on the mailing list. When accounting for the long-term effects of unsubscriptions on giving, we find that the net effect for the charity of sending a reminder is just 1 DKK (\$0.18) per potential donor.

Our paper adds to several strands of literature. First, it presents evidence on the hidden costs of nudging that have been mostly neglected until now. Previous studies on reminders do not distinguish between an effect on behavior and an effect on welfare (Macharia et al., 1992; Vervloet et al., 2012; Karlan et al., 2012; Gilbert and Zivin, 2014; Huck and Rasul, 2010; Sonntag and Zizzo, 2015; Calzolari and Nardotto, 2014; Allcott and Rogers, 2014). However, a small number of studies have looked at the welfare effects of other choice enhancing policies (Carroll et al., 2009; Bernheim et al., 2015; Bhattacharya et al., 2015; Murooka and Schwarz, 2016). Allcott and Kessler (2015) present an experimental design to evaluate the welfare effects of social comparisons in reducing energy consumption by eliciting the willingness to pay for the nudge. They show that ignoring “non-energy costs” such as psychological costs (like guilt or shame) and time costs for turning off lights and adjusting thermostats overstates the welfare gain of the nudge by a factor of five. Chesterley (2015) makes a related point in his paper analyzing the theoretical drawbacks of default choices. Handel (2013) shows that nudges to overcome inertia in health insurance markets could exacerbate adverse selection and thus lead to an overall welfare loss through the nudge. In contrast to the other studies, we estimate welfare effects both for the one nudging and for the one being nudged.

Second, our paper provides new insights on social pressure costs which have been discussed in the “avoiding-the-ask” literature (Dana et al., 2007; Andreoni et al., 2011; DellaVigna et al., 2012; Knutsson et al., 2013; Cain et al., 2014; Trachtman et al., 2015). By observing the decisions of opting in or out of a possibility to donate, we observe the costs of the nudge. If the social pressure/annoyance costs are too high in our experiment, people unsubscribe from the mailing list, while they avoid opening the door or tick the opt-out box in DellaVigna et al. (2012).

Third, the model we propose in this paper is an addition to the theoretical literature on the effect of reminders on attention (Karlan et al., 2012; Taubinsky, 2013; Calzolari and Nardotto, 2014). These papers assume that individuals have a decaying attention function, which can be reset to full attention through the use of a reminder. Although all authors agree that there seems to be a natural upper limit on the number of reminders that should be sent out, their models nevertheless predict an increasing utility in the number of reminders.

In terms of its methodology, our paper is part of a growing literature on structural behavioral economics which estimate behavioral models using non-experimental (Conlin et al., 2007; Laibson et al., 2007) or experimental field data (DellaVigna et al., 2012). A close link between the theoretical model and the field experiment is advocated in this literature (Card et al., 2011; Harrison,

2014).

The remainder of the article proceeds as follows. Section 2 presents our model, while section 3 introduces the design of the two experiments and ties testable predictions derived from our model to their treatments. Section 4 describes our sample and the implementation of our experiment. Section 5 presents the reduced-form results from our experiments, and section 6 estimates the structural parameters of our model. Section 7 presents our welfare analysis and section 8 concludes.

2 Model

Building on the work by Andreoni (1989, 1990), we present a dynamic T -period model of giving and unsubscription behavior which includes a fixed cost of each solicitation to the potential donor. The potential donor chooses both whether to give and whether to unsubscribe.¹

2.1 The general setup

We consider a repeated interaction between a charity and a warm-list donor who is asked to give via e-mails. We refer to the potential donor simply as “the donor” and to the solicitations as “the messages”. In every period $t \in \{1, 2, \dots, T\}$, the donor must decide if he wants to donate and if so, how much. In addition, whenever he receives a message, he decides if he wants to unsubscribe from future messages sent by the charity.

We assume that the donor receives warm-glow utility from every donation $g_t \geq 0$ to the charity. We denote the warm-glow utility from giving by $v(g_t)$ where $v'(\cdot) > 0$, $v''(\cdot) < 0$, and $\lim_{g_t \rightarrow \infty} v'(g_t) = 0$.² We model the cost of giving by the function $c(\cdot)$ where $c'(\cdot) > 0$ and assume that this captures all costs associated with giving, including the reduction in consumption utility, transaction costs, and opportunity costs. The net donation utility from giving g_t is therefore

$$d(g_t, a_t) = a_t v(g_t) - c(g_t)$$

¹The technical details, including proofs, are provided in the Supplementary Appendix.

²Note that although we refer to it as warm-glow utility, $v(\cdot)$ could also capture prestige or utility from conforming to social norms, and the model could easily be adapted to include pure altruism.

where a_t is the weight on warm-glow utility. We further assume that $d''_{gg}(g_t, a_t) < 0$, $d(0, a_t) = 0$, and $d(g_t, a_t) \in L^1$, i.e., the integral of the absolute value of $d(g_t, a_t)$ is finite. The law of motion for a_t is given by an AR(1) process³

$$a_t = \mu + \rho a_{t-1} + \varepsilon_t$$

where $\varepsilon_t \sim IID(0, \sigma^2)$ on a finite support $[-M, M]$, i.e., $M < \infty$.

To capture the effect of reminders, we assume that the donor has limited attention and therefore only remembers the donation problem with probability $\theta \in [0, 1)$ in every period. If the donor is attentive and remembers the donation decision, he gives an amount $g_t \geq 0$ to the charity. On the other hand, if he is inattentive and forgets about the donation decision, then $g_t = 0$. Similar to the inattention models by Karlan et al. (2012) and Taubinsky (2013), we assume that the donor is sophisticated and therefore aware of his inattention.

We assume that any message from the charity serves as a reminder of the donation problem. We let p_t denote the probability that the charity sends a message in period t . The donor receives the message if he has not unsubscribed in any of the previous periods. If the donor is subscribed to messages in period t and the charity sends a message, then the donor always recalls the donation problem, otherwise the donation problem is only remembered with probability θ .⁴ Hence, subscribing to the mailing list at the beginning of period t , increases the probability that the donor remembers the donation problem.

We let Λ denote a cost to the donor of receiving a message from the charity. This cost can be thought of as an effort cost of looking at the message or a moral cost of feeling guilty for having to be reminded. We refer to this cost as an “annoyance cost”, which for simplicity is assumed to be constant. We also assume that any type of message generates the same fixed cost, i.e., original solicitations and reminders induce the same cost.

If the donor receives a message in period t , he also has the option to unsubscribe $u_t = 1$ or not $u_t = 0$ from the mailing list. The decision to unsubscribe is considered irreversible and eliminates all future messages from the charity, i.e., $u_{t+k} = 1$ if $u_t = 1$ for all $k \in \{1, 2, \dots, T - t\}$. It follows

³Note that the AR(1) process introduces time-variation in the weight on warm-glow utility and that a deterministic process for a_t would lead to a static problem where the donor either unsubscribes in period $t = 1$ or never unsubscribes.

⁴We note that θ can capture both natural recall and cues other than direct messages, e.g., general advertisements.

that if $p_t(1 - u_{t-1}) = 1$, the donor is subscribed at the beginning of period t and receives a message from the charity.

Under these assumptions, the donor's inter-temporal optimization problem in period t is

$$\max_{g_t, u_t} E \left[\sum_{\tau=0}^{T-t} \delta^\tau \left[p_{t+\tau}(1 - u_{t-1+\tau})(d(g_{t+\tau}, a_{t+\tau}) - \Lambda) + (1 - p_{t+\tau}(1 - u_{t-1+\tau}))\theta d(g_{t+\tau}, a_{t+\tau}) \right] \middle| \Omega_t \right] \quad (1)$$

where $0 < \delta < 1$ is the inter-temporal discount factor and $E[\cdot | \Omega_t]$ denotes the expectation given period t information. The information set Ω_t includes $\{a_\tau\}_{\tau=1}^t$, meaning that the donor knows the weight he assigned to warm-glow utility in current and past periods.

A few remarks regarding the dynamic structure of the model are in place. First, we assume that donors have a finite horizon T to capture that people are unlikely to plan years ahead when it comes to charitable giving. In addition, a fixed horizon T allow us to investigate the effect of varying the horizon. Second, we do not assume an inter-temporal budget constraint for the maximization problem. This simplifying assumption is made for tractability and because in the case of charitable giving it seems unlikely that the inter-temporal budget constraint would be binding. Prediction 5 below follows directly from this assumption and therefore allow us to test this feature of the model.

When solving the model we assume that the donor is rational in the sense that he knows his preferences, the timing of events, how he will respond to messages in future periods, and forms rational expectations regarding the charity's reminder strategy, i.e., $\{p_t\}_{t=1}^T$. If a donor has not unsubscribed in period t but does not receive a message because $p_t = 0$, then he has no opportunity to unsubscribe, and we therefore let $u_t = 0$. In addition, it is assumed that the donor does not unsubscribe when he is indifferent between doing so or not.⁵ Similarly, we assume that the donor does not give anything if he is indifferent between doing so or not.

2.2 Giving and unsubscribing

The optimal donation and unsubscription decisions are obtained by backwards induction, and both have classic threshold properties. As illustrated in Figure 1, a donor with a sufficiently low real-

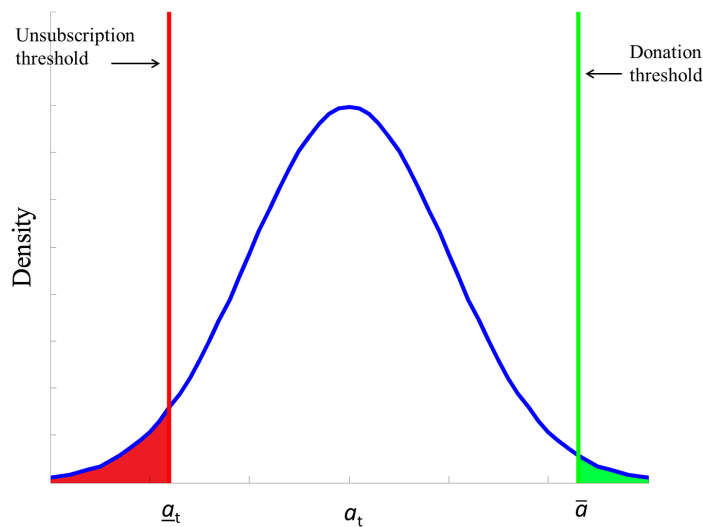
⁵One can think of this assumption as capturing some tiny cost of pressing the unsubscribe button that tips the balance in favor of not unsubscribing. Alternatively, it can be interpreted as a small default bias.

ization of a_t unsubscribes, while a donor with a high realization of a_t makes a donation. Lemma 1 gives the threshold property for the optimal donation choice.

Lemma 1. *Conditional on remembering, the optimal donation decision $g^*(a_t)$ is weakly increasing in a_t and has the threshold property: (i) $g^*(a_t) = 0$ for $a_t \leq \bar{a} \equiv \frac{c'(0)}{v'(0)}$; and (ii) $g^*(a_t) > 0$ for $a_t > \bar{a}$.*

Lemma 1 follows directly from the maximization problem in Equation (1), and the assumption that giving in period t only affects utility in this period. Hence, conditional on remembering the donation problem, the donor only takes into account his current generosity, as captured by a_t , when choosing whether to make a positive donation. The amount donated is weakly increasing in a_t .⁶ Notice that the threshold \bar{a} is constant and that time variation in giving solely originates from variability in a_t .

Figure 1: Optimal unsubscription and donation thresholds



Notes: The optimal donation and unsubscription thresholds when ε_t follows a truncated normal distribution.

Lemma 2 gives the corresponding threshold property for the optimal unsubscription choice. Recall that M is the upper bound on the support for the innovation in the $AR(1)$ -process for a_t .

⁶This is similar to the predictions of usual static models of giving in Andreoni (1989, 1990) and DellaVigna et al. (2012).

Lemma 2. *Conditional on receiving a message from the charity in period t and for M sufficiently large, the optimal unsubscription decision $u_t^*(a_t)$ has the threshold property: (i) $u_t^*(a_t) = 1$ for $a_t < \underline{a}_t$; and (ii) $u_t^*(a_t) = 0$ for $a_t \geq \underline{a}_t$. The threshold \underline{a}_t is increasing in t .*

Lemma 2 follows from exploiting the finding in Lemma 1 that g^* only depends on a_t and not on current or future unsubscription or donation choices. This property of our model eases tractability and implies that the model solution has a sequential structure: the donor simply conditions on the optimal donation rule when making his unsubscription decision. More formally, for a given a_t , he first obtains $\{g_\tau^*\}_{\tau=t}^T$ and then computes $\{u_\tau^*\}_{\tau=t}^T$ backwards. The unsubscription problem is then reduced to an optimal stopping problem with one state variable a_t and one control variable $u_t \in A = \{0, 1\}$. Its value function is given by a standard Bellman equation, and it is well-known that the solution has a threshold property (see for instance Rust (1987)).⁷

The lower threshold \underline{a}_t can be interpreted as an optimal unsubscription boundary similar to the optimal exercise boundary for American options.⁸ In our context, the donor unsubscribes if he expects future annoyance costs to be larger than the warm-glow utility foregone by not being reminded in the future. This latter effect is referred to as “the option value” of subscribing. The boundary \underline{a}_t increases with time as the value of subscribing decreases over time given a finite horizon T .

The threshold property in Lemma 2 implies that a donor unsubscribes from future messages if he gets a message and his realization of a_t is sufficiently low. Combining Lemma 1 and Lemma 2, we obtain the following proposition for the effect of reminders:

Proposition 1. *A reminder increases both the unconditional probability that the donor makes a donation and the unconditional probability that he unsubscribes.*

We note that current utility is unaffected by the unsubscription choice in this period because the current annoyance cost cannot be avoided. Hence, the current unsubscription choice only affects the utility in future periods. We summarize these observations in the following proposition which ties the current unsubscription decision to expectations about the future:

⁷Formal proofs verifying these results in our setting are deferred to the Supplementary Appendix. The technical requirement that the bound M (on the support for ε_t in the AR(1) process) must be sufficiently large is not particularly restrictive as M can be chosen arbitrarily large.

⁸For instance, the holder of an American put option exercises his right to sell the option if its value falls below a critical value, which is referred to as the optimal exercise boundary (Kim, 1990; Jacka, 1991; Carr et al., 1992, among others).

Proposition 2. *The unsubscription choice of the donor depends on the option value of subscribing and expected future annoyance costs.*

3 Experimental design and testable predictions

To test our model, and Proposition 1 and 2 in particular, we design two field experiments carried out via e-mail. The following sub-sections describe these experiments, their treatments, and derive testable predictions which hold for *all* specifications of warm-glow utility and costs given the stated assumptions.

3.1 Experiment I: A targeted reminder

Experiment I tests Proposition 1 that reminders increase donations at the expense of more unsubscriptions. The experiment is carried out in a setting with infrequent e-mail communication from the charity to donors. Potential donors are randomized into two treatments:⁹

- **Control I (CI):** A solicitation e-mail presents the cause and informs that for every person who donates within the next 10 days an anonymous donor will donate an additional 10 DKK (approx. 1.8 USD).
- **Targeted Reminder (TR):** In addition to the first e-mail an unannounced targeted reminder is sent seven days later to anyone who has not donated or unsubscribed within the first week. The reminder contains no new information.

Given the specific treatments, Proposition 1 can be translated into testable predictions. In both treatments, we send an initial donation request in period $t = 1$. In addition, potential donors in the Targeted Reminder treatment receive an unannounced targeted reminder in period $t = 2$ if they do not give or unsubscribe in response to the initial message. Let u_t^j and g_t^j denote the unsubscription and donation decision, respectively, in period t for individuals in treatment j where $j \in \{CI; TR\}$.

Prediction 1. *The unconditional probability of giving is higher in the Targeted Reminder treatment than in the Control I treatment: $P(g_1^{TR} + g_2^{TR} > 0) > P(g_1^{CI} + g_2^{CI} > 0)$. In particular, the*

⁹A number of other cross-randomized treatments were implemented in parallel and are described in detail in Damgaard and Gravert (2014). This paper focuses on the two treatments that allow us to isolate the effect of the targeted reminder.

unconditional probability of giving in period $t = 2$ is larger in the Targeted Reminder treatment than in the Control I treatment: $P(g_2^{TR} > 0) > P(g_2^{CI} > 0)$.

Donors in the two treatments are equally likely to give in period $t = 1$. However, donors in Control I do not receive the targeted reminder and are inattentive to the donation problem in period $t = 2$ with probability θ . They may therefore fail to give in period $t = 2$ even if their realization of a_2 is above the upper threshold \bar{a} . Potential donors in the Targeted Reminder treatment remember the donation problem in period $t = 2$ with certainty and are therefore more likely to donate in period $t = 2$. Taking the two periods together, the probability of giving is therefore greater in the Targeted Reminder treatment than in Control I.¹⁰

Next we consider the probability of unsubscribing.

Prediction 2. *The unconditional probability of unsubscribing is higher in the Targeted Reminder treatment than in the Control I treatment: $P(u_2^{TR} = 1) > P(u_2^{CI} = 1)$.*

Donors in the Targeted Reminder get two messages that they can unsubscribe from, while those in Control I only have one possibility to unsubscribe.

3.2 Experiment II: Changing the option value of subscribing

Experiment II tests Proposition 2 that unsubscription choices are affected by the option value of subscribing and beliefs about future annoyance cost. We therefore consider treatments that either i) change the option value of subscribing by announcing a future “matching” opportunity, or ii) change the future expected annoyance costs by announcing a temporary reduction in the frequency of messages. The experiment is carried out in a setting with e-mails from the charity to donors approximately once a month.

Potential donors are randomized equally into the following three treatments:¹¹

- **Control II (CII):** A solicitation e-mail informs of the cause and contains the information that subscribers usually receive one e-mail a month from the charity.
- **Future Benefit (FB):** The same solicitation e-mail as in Control II is used with the information that subscribers usually receive one e-mail a month plus an announcement that in the

¹⁰This prediction of the model is similar to the prediction in Huck and Rasul (2010).

¹¹We cross-randomize the receivers who participated in Experiment I and the new subscribers into the treatments of Experiment II to avoid any confounds with the first experiment.

next e-mail an anonymous donor will donate a healthy meal to a poor child for every person on the mailing list who donates in response to the second e-mail.

- **Low Frequency (LF):** The same solicitation e-mail as in Control II is used with the information that subscribers usually receive one e-mail a month plus an announcement that in the coming three months subscribers will only receive one e-mail from the charity.

To derive testable predictions related to Proposition 2, we let one period correspond to one month. Hence, $p_t = 1$ for all t in the Control II and Future Benefit treatments, as the charity sends a message every month and potential donors are made aware of this. In the Future Benefit treatment potential donors are told in period $t = 1$ that a lead donor will give a “match” worth $m > 0$ for every person on the mailing list who donates at least $X > 0$ in period $t = 2$. Assume that the donors get warm-glow utility from the sponsored amount m , then donors in the Future Benefit treatment on the mailing list in period $t = 2$ (i.e. $u_1^{FB} = 0$) get donation utility

$$d_m(g_2, a_2) = \begin{cases} a_2 v(g_2 + m) - c(g_2) & \text{if } g_2 \geq X \\ a_2 v(g_2) - c(g_2) & \text{otherwise.} \end{cases}$$

Donors in the Future Benefit treatment who are not on the mailing list at time $t = 2$ get the standard donation utility, i.e., $d(g_2, a_2) = a_2 v(g_2) - c(g_2)$. Assume that donors in the Control II treatment believe that $m = 0$, and assume that the utility in all other periods is unaffected by the announcement of the match. This leads to a prediction about the relative size of unsubscription rates in the Future Benefit and Control II treatments.

Prediction 3. *The unconditional probability of unsubscribing in period $t = 1$ is lower in the Future Benefit treatment than in the Control II treatment : $P(u_1^{FB} = 1) < P(u_1^{CII} = 1)$.*

To understand this result, first note that the match $m > 0$ reduces the cost of achieving a certain level of warm-glow utility for donations above the threshold X and thus $d_m(g_2^*, a_2) \geq d(g_2^*, a_2)$. Hence for a given value of a_1 , the expected option value of remaining subscribed (at least until the next period) is greater in the Future Benefit treatment than in the Control II treatment. The model therefore predicts that the unsubscription rate is smaller in the Future Benefit treatment than in the Control II treatment.

To derive a similar prediction for the Low Frequency and Control II treatments, note that donors in the Low Frequency treatment in period $t = 1$ are told that they will only receive one message in

the next three months, i.e., $p_2 = p_3 = p_4 = \frac{1}{3}$ and $p_j = 1$ for all $j > 4$. This leads to a prediction about the size of the unconditional unsubscription rates.

Prediction 4. *For $P(a_t > \bar{a})$ sufficiently low, the unconditional probability of unsubscribing in period $t = 1$ is lower in the Low Frequency treatment than in the Control II treatment: $P(u_1^{LF} = 1) < P(u_1^{CII} = 1)$.*

The intuition behind this result is as follows. The Low Frequency treatment implies fewer messages that impose annoyance costs on potential donors. At the same time potential donors are less likely to remember to give in cases where giving is optimal, and this could make it less valuable to remain on the list. It turns out that if there is a relatively small probability of giving, i.e., $P(a_t > \bar{a})$ sufficiently low, then the effect of lower annoyance outweighs the effect of foregoing opportunities to donate. This therefore implies that the unconditional probability of unsubscribing is lower in the Low Frequency treatment than in the Control II treatment.¹²

In terms of giving behavior, the model predicts no differences across treatments:

Prediction 5. *The unconditional probability of giving in period $t = 1$ is the same in the Control II treatment, the Low Frequency treatment, and the Future Benefit treatment: $P(g_1^{CII} > 0) = P(g_1^{LF} > 0) = P(g_1^{FB} > 0)$.*

This prediction arises directly from the assumption that giving is not constrained by an inter-temporal budget constraint. Effectively, Prediction 5 is a test of the validity of this assumption.¹³

4 Sample and implementation

We collaborated with the Danish charity DanChurchAid (DCA) to run the field experiments in the summers of 2013 and 2015. DCA is one of the largest NGO's in Denmark with a total revenue of 0.57 billion DKK in 2013 (DanChurch Aid, 2013). The total annual revenue of charities in Den-

¹²In a setting as ours with no social pressure costs, it seems reasonable to assume that $P(a_t > \bar{a})$ is small. In Control I we find very little giving without asking. This is similar to the results reported by DellaVigna et al. (2012) who find virtually no giving via regular mail or internet but only giving through door-to-door solicitation.

¹³Intuitively, it implies that there is no inter-temporal substitution of giving as only the current realization of a_t matters for the decision to give.

mark is estimated to be about 2 billion DKK.¹⁴ DCA mostly implements and supports emergency and development programs in Asia, Africa, the Middle East, and Central America.

Our samples for the two experiments consist of warm-list donors who have provided their e-mail address to the charity. The mailing list is constantly updated as new subscribers are added and others unsubscribe or close their e-mail accounts. A total of 11,324 individuals (roughly one third of the combined sample) participated in both experiments. Our samples do not include regular donors with payments to the charity setup as a monthly Direct Debit at the time of the experiments because the automatic nature of these payments alter the attention considerations. E-mail communication from the charity was relatively uncommon prior to the first experiment and varied depending on which campaigns donors had previously responded to. However, at the time of the second experiment, donors on the mailing list had received e-mail messages from the charity approximately every month for the past year. In addition to e-mails, the charity uses several other communication channels to reach potential donors, including mass media, social media, regular door-to-door solicitations, and text messages solicitations. DCA also runs 125 charity shops across Denmark, and it has partnered with an electricity provider to offer people the opportunity to donate via their electricity bill. All donations are tax deductible, which is stated in all correspondence.

4.1 Implementation of the experiments

The initial e-mail in Experiment I was sent on the 28th of May 2013, and the reminder was sent on the 4th of June 2013. Our sample for the first experiment consisted of 17,391 donors, and approximately half the sample was randomly allocated to each of the two treatments (Targeted Reminder and Control I). Personal characteristics are similar across the two treatments as shown in Table 1. The style of the e-mail was similar to the style of other communication sent by the charity, and the e-mail solicited money for poor children in Africa (see Figure A3 for a screenshot).¹⁵

For Experiment II, 43,591 donors received a solicitation e-mail on the 9th of July 2015. The e-mail was in the style of regular solicitations by the charity and announced the possibility of supporting the opening of a store selling surplus food in order to reduce food waste and raise money for the charity (see Figure A4 for a screenshot). People were asked to donate money in steps of 100 DKK, which constituted the “price” of a “share” in the store, but the shares did not

¹⁴Deloitte and the Danish Fundraising Association (ISOBRO) estimate that ISOBRO members had a combined revenue of 1.8 billion DKK and accounted for more than 75% of the market in 2013 (ISOBRO and Deloitte, 2014).

¹⁵Translations of the experimental material are in the Supplementary Appendix.

entitle them to any ownership or rights regarding the store, i.e., it was a pure donation. Donors did however receive a physical printout of a share in the mail for every 100 DKK donation they made.¹⁶ To avoid self-selection into opening the e-mail, all three treatments had the same subject line “Stop Food Waste”.

A second e-mail was sent out a month after Experiment II to measure the medium term effects of the intervention and provide the matching opportunity announced in the Future Benefit treatment. More information on this e-mail and the effect it had can be found in Appendix C.

We obtained a good balance across the treatments in Experiment II, as shown in Table 1. Given the natural development in the e-mail list of the charity, some of the summary statistics have changed between the first and the second experiment. The average age is lower (38 versus 46 years), and the average amount donated at the last donation through any channel has decreased from around 300 DKK to around 190 DKK. Other characteristics are very similar to those of the first experiment.

4.1.1 The unsubscribe link and landing page

It was possible to unsubscribe from the mailing list by clicking a button at the bottom of every e-mail. The design and visibility of the button was identical in all e-mails. If donors clicked on the unsubscribe button, they were directed to a website hosted by the charity, a so-called landing page. In Experiment I the landing page would prompt donors to confirm the unsubscription. In Experiment II we used the landing page to gather survey information about *why* donors unsubscribe, thus complementing the experimental treatments. The landing page therefore presented unsubscribers with five radio buttons; four possible reasons for unsubscribing and an Other choice, allowing them to specify a reason. Two of the stated reasons were generic and allowed donors to express a general lack of interest in the charity and newsletter (“I no longer want to give to DCA” and “I don’t find the content of the newsletter interesting”). The other two reasons provided information about the role of annoyance (“DCA sends me too many e-mails”) and perceived pushiness of the charity (“I don’t like to be asked directly to donate to DCA”) in the unsubscription decision. Unsubscribers were asked to choose one of the five options provided and confirm the unsubscription.

¹⁶The physical “food share” was an illustrated sheet of paper. If donors were not interested in receiving the physical share, they could opt out and make a “regular” donation to the project. In our sample, 78 percent of the donors received the physical share. None of the explanatory variables we have in our data significantly explain opting out of receiving a share in a probit regression (results available on request).

Table 1: Summary statistics and covariates balance

	Experiment I		Experiment II		
	Control I	Targeted Reminder	Control II	Low Frequency	Future Benefit
Female (share)	0.62 (0.49)	0.63 (0.48)	0.63 (0.48)	0.64 (0.48)	0.63 (0.48)
Age (years)	46 (15)	46 (15)	38 (15)	38 (15)	38 (15)
City (share)	0.33 (0.47)	0.33 (0.47)	0.35 (0.48)	0.36 (0.48)	0.35 (0.48)
Amount donated last time (DKK)	300 (622)	313 (553)	191 (408)	194 (419)	192 (516)
Number of months since last donation	35 (19)	35 (19)	32 (22)	32 (22)	31 (22)
Number of months on e-mail list	1 (-)	1 (-)	24 (5.5)	24 (5.5)	24 (5.5)
Observations	8,692	8,699	14,536	14,527	14,528

Notes: The table reports means and standard deviations (in brackets). The variable city is a dummy for the 10 biggest cities in Denmark. For Experiment I (Experiment II) information on city is available for 99% (87%) of the sample, gender for 83% (89%), age for 41% (70%), and past donations for 88% (76%) of the sample. The number of months on the e-mail list was at most 27 months in Experiment II. By definition it was equal to one month in Experiment I.

On the next page the unsubscription was confirmed, and a link to the general homepage of the charity was provided.

5 Reduced-form results

We first present the results on giving and unsubscriptions from Experiment I before presenting the results from Experiment II. An overview of the response rates and total observations for the two experiments is provided in Table 2. The experiments have similar donation rates, but there is a relatively large drop in the unsubscription rate from Experiment I to Experiment II which we discuss further in section 5.2.1. The average amount donated is similar across the two experiments although the causes supported by the two experiments are different.

Table 2: Experiment I and II: Results statistics

	Experiment I			Experiment II			
	All	Control I	Targeted Reminder	All	Control II	Low Frequency	Future Benefit
Responses (in %)							
Percentage who gave	0.44 (6.6)	0.35 (5.9)	0.53 (7.3)	0.66 (8.1)	0.65 (8.0)	0.67 (8.2)	0.65 (8.0)
Percentage who unsubscribed	2.90 (16.8)	2.14 (14.5)	3.67 (18.8)	0.38 (6.9)	0.49 (7.0)	0.30 (5.5)	0.36 (5.9)
Observations (N)							
Full sample	17,391	8,692	8,699	43,489	14,501	14,494	14,494
Number of people who gave	76	30	46	285	94	97	94
Number of unsubscribers	504	186	318	167	71	44	52

Notes: The table provides means and standard deviations (in brackets).

5.1 Experiment I: The effect of a targeted reminder

Our model provides a number of predictions regarding giving, timing of giving, and unsubscription behavior in Experiment I. We discuss the evidence for each of these predictions in turn.

5.1.1 The reminder increases the number of donations

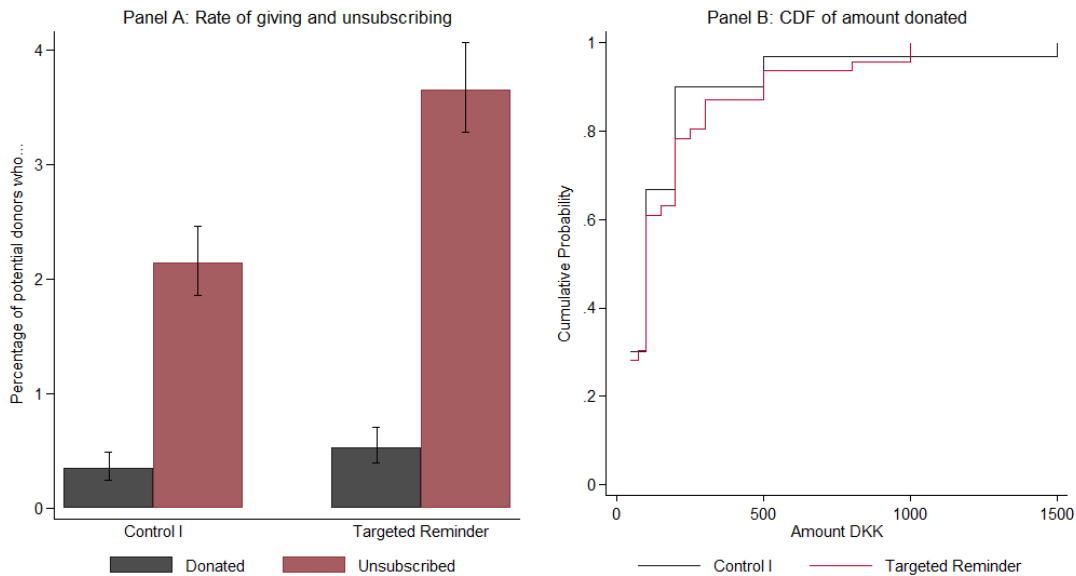
Figure 2 shows that the share of donors with positive donations in the Targeted Reminder treatment is larger than that in the Control I treatment (0.53% and 0.35%, respectively), which is in line with Prediction 1. This is a significant increase of about two-thirds (p-value = 0.066). In addition, Table 3 provides the results of probit regressions on the likelihood of donating. The treatment effect is similar in sign and magnitude when including individual specific controls but insignificant. We find that age has a significant positive effect on donating, while the probability of donating is negatively related to the amount of time that has passed since the last donation was made. We do not see any significant effect of gender, place of living, or the most recent amount donated on the probability of donating. Thus, we replicate the findings of Huck and Rasul (2010) that reminders can increase donations on the extensive margin. When it comes to the intensive margin, we see a slight increase in the amount donated, conditional on donating, for people in the Targeted Reminder group compared to the Control I group as shown in Figure 2. However, this increase is not significant and does not hold in a regression analysis of the amount donated unconditional on donating (see Table A1 in the appendix).

Table 3: Donation decisions in Experiment I and II

	Experiment I		Experiment II	
	(1)	(2)	(3)	(4)
Targeted Reminder	0.00184* (0.00100)	0.00269 (0.00195)		
Low Frequency			0.00021 (0.00095)	0.00006 (0.00091)
Future Benefit			0.00000 (0.00095)	-0.00006 (0.0009)
Female		0.00131 (0.00193)		0.00198*** (0.00074)
Age		0.00015** (0.00007)		0.00026*** (0.00002)
City		0.00058 (0.00214)		0.00319*** (0.00093)
Months since last donated		-0.00027*** (0.00006)		-0.00005** (0.00002)
Amount last donated		-0.00000 (0.00000)		0.00000*** (0.00005)
Months on e-mail list				-0.00018*** (0.00000)
Observations	17,391	6,448	43,592	27,220
Pseudo R^2	0.004	0.046	0.000	0.079

Notes: The table provides the marginal effects and standard errors in brackets of probit regressions on the binary donation choice. The variables Targeted Reminder, Low Frequency, and Future Benefit are dummy variables that are evaluated in comparison to their respective control groups (control dummies are set to zero). Female and City are dummy variables. Months since last donated and Amount last donated correspond to the last donation prior to the respective experiment through any channel. Months on e-mail list is set at one month for everyone in the first experiment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

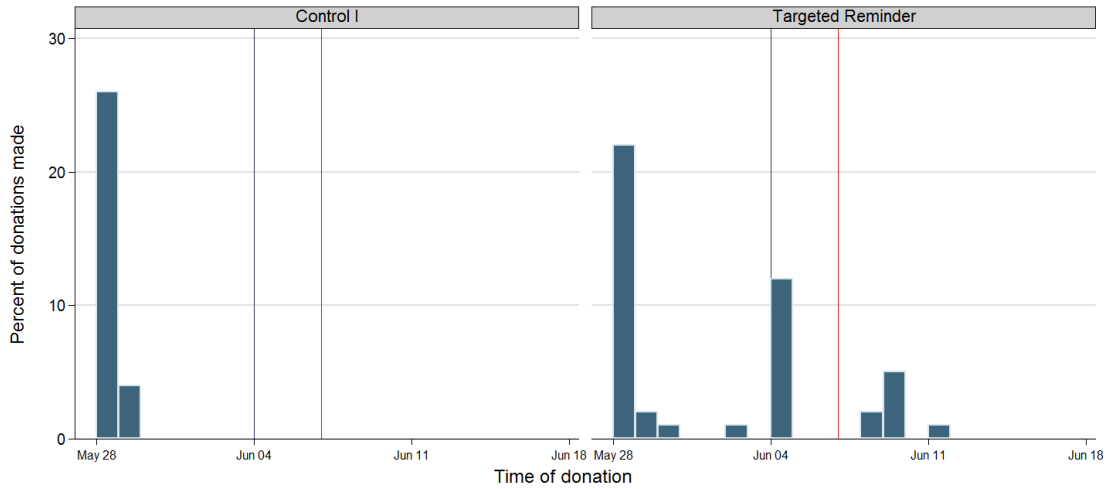
Figure 2: Giving and unsubscription behavior in Experiment I



Notes: Panel A illustrates the rate of giving and unsubsribing with confidence intervals. Panel B shows the cumulative distribution function (CDF) of the amount donated conditional on giving. There are 30 and 46 donors in Control I and Targeted Reminder, respectively. The corresponding number of unsubscribers is 186 and 318, respectively. Difference in rate of giving is significant at 10% level (Pearson $\chi^2(1) = 3.3703$, p-value = 0.066, Fisher's exact = 0.084) and the difference in unsubscription rate is significant at 1% level (Pearson $\chi^2(1) = 35.4939$, p-value = 0.000, Fisher's exact = 0.000). The differences in distribution of amounts between the treatments are not significant using a Mann-Whitney test (p-values > 0.62), and neither do we find a significant difference in a two-sided two-sample t-test or a Kolmogorov-Smirnov test.

We find further support for Prediction 1 when looking at the timing of the donations. Prediction 1 says that the probability of donating around the time of the reminder is higher in the Targeted Reminder group than in the Control I group. Figure 3 shows that the donations are made either on the day or the day after the initial solicitation mail or the reminder, and we only see donations around the time of the reminder in the Targeted Reminder group. Hence, donations are *not* made close to the deadline. The results suggest that receivers have a very low rate of natural recall and are unlikely to make a donation without being reminded.

Figure 3: Timing of giving in Experiment I



Notes: The figure shows the percentage of donations made on each day in the Control I and Targeted Reminder groups. The initial e-mail was sent on May 28th, and the reminder was sent on June 4th. The deadline for donating was June 7th. Data for timing of unsubscriptions is not available for Experiment I.

5.1.2 The reminder increases the number of unsubscriptions

Figure 2 documents a large treatment difference in unsubscription behavior.¹⁷ The Targeted Reminder is associated with a higher unsubscription rate of 3.7% compared to an unsubscription rate of 2.1% in Control I. This is a difference of about 76%, it is highly significant (p-value = 0.000), and it is in line with Prediction 2. In the regression analysis provided in Table 4, we find that this effect is robust to the inclusion control variables. Including all controls, the reminder increases the likelihood of unsubscribing by 1.1 percentage points compared to Control I. We further find some evidence that women are significantly less likely to unsubscribe from the mailing list, as are individuals who are older. The place of living and the amount last donated have no effect on unsubscriptions.

5.2 Experiment II: Effects of changing the option value

We now present the results from Experiment II which change the option value of subscribing in two treatments. Here we test whether people account for the option value of subscribing when making their unsubscription decision as predicted by our model.

¹⁷Data on unsubscriptions is only available for the treatment period, and we only have information on unsubscriptions through the links in the e-mails sent out as part of this experiment.

Table 4: Unsubscriptions in Experiment I and II

	Experiment I		Experiment II	
	(1)	(2)	(3)	(4)
Targeted Reminder	0.01516*** (0.00254)	0.01114*** (0.00398)		
Low Frequency			-0.00186** (0.00074)	-0.00192** (0.00081)
Future Benefit			-0.00131* (0.00076)	-0.00039 (0.00091)
Female		-0.01804*** (0.00444)		0.00027 (0.00083)
Age		-0.00035** (0.00014)		-0.00003 (0.00003)
City		0.00645 (0.00445)		-0.00001 (0.00086)
Months since last donated		-0.00022* (0.00011)		-0.00001 (0.00002)
Amount last donated		0.00000 (0.00000)		-0.00000 (0.00000)
Months on e-mail list				-0.00015** (0.00007)
Observations	17391	6448	43489	27053
Pseudo R^2	0.008	0.023	0.003	0.014

Notes: The table shows the marginal effects and standard errors in brackets of probit regressions on unsubscribing. The variables Targeted Reminder, Low Frequency, and Future Benefit are dummy variables that are evaluated in comparison to their respective control groups (controls are set to zero). Female and City are dummy variables. Months since last donated and Amount last donated correspond to the last donation prior to the respective experiment through any channel. Months on e-mail list is set at one month for everyone in the first experiment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.2.1 A lower frequency of messages reduces the unsubscription rate

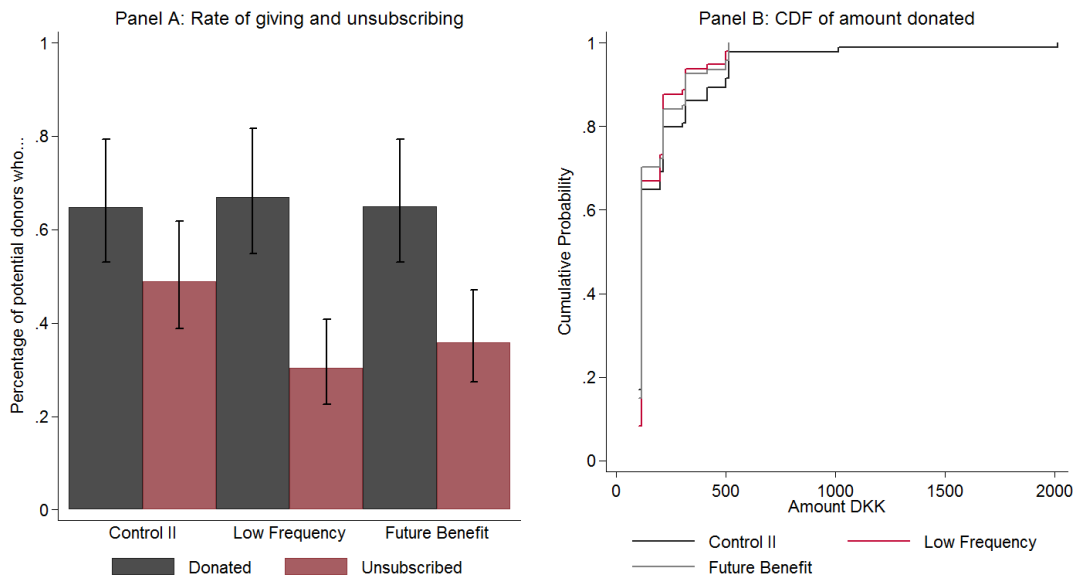
In line with Prediction 4 from our model, the Low Frequency treatment reduces the unsubscription rate from 0.49% to 0.30% (see Figure 4).¹⁸ That is a reduction of 39%. The probit regressions in Table 4 show that the announcement of the reduced frequency has a significant effect on the unsubscription rate. Including all controls, individuals in the Low Frequency treatment are 0.17-0.20 percentage points less likely to unsubscribe than donors in the Control II treatment. The coefficient of the Future Benefit treatment goes in the direction implied by the model, i.e., Prediction 3. We find a marginally significant effect on the unsubscriptions compared to the Control II treatment, but this effect is not robust to the inclusion of controls. Contrary to Experiment I, we find no effect of gender, age, or place of living on unsubscribing. This could be due to the sample size, but also to the particular good that was solicited for. The most recent donation has no measurable effect on unsubscribing, but the seniority of a donor reduces the probability of unsubscribing significantly, although the effect is small compared to the effect of the Low Frequency treatment dummy.

Surprisingly, the unsubscription rate is far lower than that in Experiment I. When we compare the unsubscription rates of our experiments with the rates the charity observed for some of their other campaigns, we find that Experiment I is at the upper range of unsubscription rates and Experiment II at the lower range. Appendix Figure A1 shows the trend in the unsubscription rate over the past two years since Experiment I. The rates have been constantly declining, with no visible difference between donation request e-mails and other newsletters. Since the e-mails of Experiment I were some of the first e-mails the donors received, individuals who dislike newsletters may have reacted to that first e-mail and left the mailing list by the time we ran Experiment II. Nevertheless, there is a constant stream of unsubscriptions every time the charity sends out an e-mail, and these subscribers are not just the most recent people joining the list (although being on the list for a longer time significantly reduces the propensity to unsubscribe). To ensure that the lower unsubscription rate in the second experiment is not explained by a more difficult unsubscription process in Experiment II due to the attached survey question, we note that the number of donors who click on the unsubscribe link (which was identical in the two experiments) in Experiment II is 222, and 167 ultimately unsubscribe. While some people seem to change their mind after clicking the link

¹⁸To measure the reaction to the reminder in a clean way, we only analyze behavior within the first three days of receiving the solicitation. This is the time frame in which most unsubscriptions are carried out, and it helps reduce the noise created by other reminders or motivations for unsubscribing such as cleaning up the inbox after summer vacation. Ideally, we would like to measure the response immediately after sending the message such that behavior (for some people) is not the result of several shocks to the weight on warm-glow utility.

to unsubscribe, this cannot explain the far lower unsubscription rate compared to Experiment I.

Figure 4: Giving and unsubscription behavior in Experiment II



Notes: Panel A illustrates the rate of giving and unsubsribing with confidence intervals. Panel B shows the cumulative distribution function (CDF) of the amount donated conditional on giving. There are 94, 97, and 94 donors in the Control II, Low Frequency, and Future Benefit treatments, respectively. The corresponding numbers of unsubscribers are 71, 44, and 52, respectively. The differences in rate of giving are not significant. The difference in unsubscription rates between Control II and Low Frequency is significant at 1% level (Pearson $\chi^2(1) = 6.35$ (p-value = 0.01)), between Control II and Future Benefit at the 10% level (Pearson $\chi^2(1) = 2.94$ (p-value=0.09)), and between Low Frequency and Future Benefit are not significant (Pearson $\chi^2(1)=0.67$ (p-value 0.41)). The differences in distribution of amounts between the treatments are not significant using a Mann-Whitney test (p-values > 0.48), nor do we find a significant difference in a two-sided two-sample t-test or a Kolmogorov-Smirnov test.

For Experiment II, we have information on the timing of the unsubscriptions (see Appendix Figure A2). Most unsubscriptions happen as an immediate reaction to the e-mail (within the first 3 days). There are no visible treatment differences between the timing of the unsubscriptions. Around 70 percent of all unsubscriptions happen on the day the e-mail is sent out.

5.2.2 The option value does not influence giving

In Table 3 we show that the treatments have no significant effect on the decision to give, which is consistent with Prediction 5. When considering the average amount donated, we find no significant effect of the treatments (see Appendix Table A1). The effects of the controls are roughly consistent with the results of Experiment I.

As in the first experiment, we find that most donations are made on the day the e-mail is sent out or the following day (see Appendix Figure A2). This shows that giving is either an immediate reaction to the solicitation or otherwise forgotten.

6 Structural estimation

Our reduced form results are consistent with all five predictions from our general model (at least in terms of directional effects) and provide suggestive evidence of a cost to potential donors for receiving reminders. To obtain a more precise estimate of this cost and to conduct a welfare analysis, we next consider a more specific version of our model which permits estimation of the structural parameters.

The version of our model considered in this section is given by $v(g_{it}) = \log(1 + g_{it})$ and $c(g_{it}) = g_{it}$.¹⁹ We allow for individual specific heterogeneity through a_{it} , where ε_{it} follows a truncated normal distribution with mean 0 and σ^2 variance on the interval $[-M; M]$ with M arbitrarily large, implying that ε_{it} effectively is normally distributed. We also let $p_t = 1$, meaning that the charity sends a message in every period, and we let a period correspond to one month.²⁰ To capture the lack of giving in some periods in the data (see below), we set the probability of remembering without being reminded to zero, i.e., $\theta = 0$. Finally, the monthly discount rate is calibrated to $\delta = 0.99835$ which corresponds to an annual real interest rate of 2%.

6.1 Solving the model

For this version of our model, we have $g_{it}^* = \operatorname{argmax}_{g_{it} \geq 0} (a_{it} \log(1 + g_{it}) - g_{it})$, which implies $g_{it}^* = a_{it} - 1$ for all $a_{it} > 1$ and $g_{it}^* = 0$ otherwise. Hence, $g_{it}^*(a_{it}) = \max(0, a_{it} - 1)$ and $\bar{a} = 1$, meaning that people with a realization of a_{it} greater than one donate a positive amount.

Given the donation rule, the unsubscription decision is derived from an optimal stopping prob-

¹⁹The assumptions of log warm-glow utility and a cost of giving proportional to the amount donated are similar to those made in DellaVigna et al. (2012).

²⁰A monthly timing is natural for the following reasons: i) at the time of Experiment II, messages were sent approximately monthly, ii) potential donors were informed of the monthly frequency in Experiment II, and iii) we observe some donors donating monthly or approximately monthly but very rarely observe more frequent donations. There are a few cases of donors making several donations on the same day. This is likely to be caused by people purchasing multiple charity “items” through the charity’s website. We treat these as one donation.

lem.²¹ Solving the optimal stopping problem is complicated by serial correlation in the unobservable state variable a_{it} , continuity of our state variable, and by the fact that the value function for the optimization problem depends non-linearly on a_{it} and its future values. This implies that we cannot obtain a closed form recursive expression for the unsubscription threshold \underline{a}_t . Instead, the solution is approximated using backwards induction, where conditional expectations are evaluated by Monte Carlo integration. Here, we use a relatively large number of draws $S = 3,000,000$ to accurately capture unsubscription and donation behavior, which are given by extremely low and high realizations of a_{it} , respectively.

6.2 Data for structural estimation

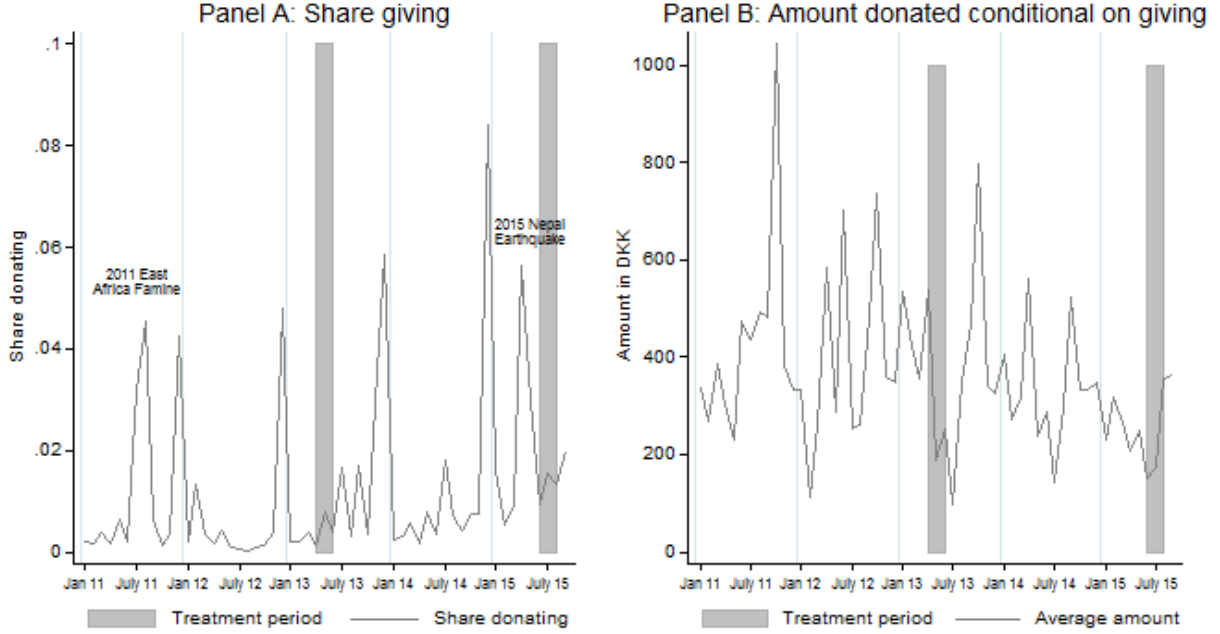
We use data from the Control II treatment in Experiment II because the frequency of messages was well-established by the time of Experiment II, and we fixed the beliefs of the donors at the correct frequency. The data is further restricted to individuals for whom historical donation data is available, giving a total of $N = 12,470$ individuals. We observe donation behavior for $\mathcal{T} = 54$ periods prior to the treatment period for Experiment II and until two periods after the treatment period.

Our data contains individual level information about amounts donated and the timing of giving. Figure 5 displays the share of people in our sample that gave a positive amount by each month, excluding donations made by Direct Debit and cash donations. All past Direct Debit donations are excluded to ensure comparability across time as our experimental sample does not include people with a Direct Debit at the time of the experiment.²² We observe large spikes in giving in December every year, which most likely capture additional giving due to Christmas and the end of the tax year. In addition, we observed spikes around the 2011 famine in East Africa and the 2015 earthquake in Nepal. Figure 5 also shows the time series for the average amount donated conditional on donating.

²¹Eckstein and Wolpin (1989), Aguirregabiria and Mira (2010), and Rust (1994) provide reviews of solution and estimation methods for dynamic stochastic discrete choice models including optimal stopping problems.

²²Cash donations for example from street solicitations or door-to-door fundraisers are not linked to donors in the database and hence cannot be included in the analysis.

Figure 5: Donation behavior over time



Notes: Based on Experiment II sample. The figure excludes payments made using Direct Debit or cash.

6.3 Estimation methodology

To present our estimation approach let $\boldsymbol{\gamma}'_1 = (\mu, \sigma, \rho)$ and $\boldsymbol{\gamma}_2 = \Lambda$ with $\boldsymbol{\gamma} = (\boldsymbol{\gamma}'_1, \boldsymbol{\gamma}_2)'$. This decomposition of our structural parameters is adopted because $\boldsymbol{\gamma}_1$ contains parameters in the process for the weight on warm-glow that can be identified solely from historical donation data independently of the donor's planning horizon T , whereas the annoyance cost in $\boldsymbol{\gamma}_2$ must be identified from unsubscription data and therefore depends on the planning horizon. To facilitate comparisons of $\boldsymbol{\gamma}_2$ for different planning horizons and hence interpretation of $\boldsymbol{\gamma}_2$, we first estimate $\boldsymbol{\gamma}_1$ using historical data prior to the treatment period, and then identify $\boldsymbol{\gamma}_2$ from the unsubscription behavior in the treatment period given our estimate of $\boldsymbol{\gamma}_1$.

6.3.1 Step 1: Estimation of $\boldsymbol{\gamma}_1$

Estimation of $\boldsymbol{\gamma}_1$ is complicated by the fact that a_{it} is unobserved, and we therefore use the method of simulated moments (MSM) following McFadden (1989).

The considered moments are: i) the probability of not giving $P(g_{it} = 0)$, ii) the probability of

giving 100 DKK or less $P(0 < g_{it} \leq 100)$, iii) the probability of giving between 100 DKK and 200 DKK $P(100 < g_{it} \leq 200)$, iv) the probability of giving between 200 DKK and 400 DKK $P(200 < g_{it} \leq 400)$, v) the probability of giving between 400 DKK and 600 DKK $P(400 < g_{it} \leq 600)$, vi) a measure of the auto-covariance in giving $E(g_{it}g_{it-1})$.²³ Our last moment is required to identify the persistence ρ in the process for a_{it} . These six empirical moments are stored in $\mathbf{m}_1(g_{it})$. The corresponding model moments $E[\mathbf{m}_1(\boldsymbol{\gamma}_1)]$ are computed by simulating donation data for \mathcal{T} periods across $\pi_1 N$ individuals, where $\pi_1 = 10$ is the scaling factor controlling the number of simulations.

The estimation is carried out using the standard procedure where the weighting matrix $\mathbf{W} = \text{diag}(\mathbf{S}_{mean}^{-1})$ in a preliminary step to obtain $\tilde{\boldsymbol{\gamma}}_1$ with \mathbf{S}_{mean} denoting the variance of $\frac{1}{N\mathcal{T}} \sum_{i=1}^N \sum_{t=1}^{\mathcal{T}} \mathbf{m}_1(g_{it})$ when re-centered around their sample means. Our final estimate $\hat{\boldsymbol{\gamma}}_1$ is then obtained using the optimal weighting matrix $\hat{\mathbf{W}} = (\hat{\mathbf{S}})^{-1}$, where $\hat{\mathbf{S}}$ denotes the variance of the moments when re-centered around $E[\mathbf{m}_1(\tilde{\boldsymbol{\gamma}}_1)]$. In both cases the variances are obtained by the Newey-West estimator using 12 lags. As shown by McFadden (1989), $\hat{\boldsymbol{\gamma}}_1$ is asymptotically normal for $N \rightarrow \infty$ with $\widehat{\text{Var}}[\hat{\boldsymbol{\gamma}}_1] = \frac{1}{NT} \left(1 + \frac{1}{\pi_1}\right) \left(\hat{\mathbf{G}}' \hat{\mathbf{W}} \hat{\mathbf{G}}\right)^{-1}$, where $\hat{\mathbf{G}} = \frac{1}{\pi_1 N \mathcal{T}} \sum_{t=1}^{\mathcal{T}} \sum_{s=1}^{\pi_1 N} \frac{\partial \mathbf{m}(\hat{\boldsymbol{\gamma}}_1, g_s)}{\partial \boldsymbol{\gamma}_1}$.

6.3.2 Step 2: Estimation of γ_2

To estimate γ_2 we condition on $\hat{\boldsymbol{\gamma}}_1$, impute starting values for a_{it} from the unconditional distribution, and determine $\hat{\gamma}_2$ using MSM with the unsubscription rate $P(u_i = 1)$ serving as our only moment. Our setting is therefore just identified, making the weighting matrix redundant. We let $m_2(u_i)$ denote the empirical moment and $E[m_2(\gamma_1, \hat{\boldsymbol{\gamma}}_1)]$ the model-implied moment. The latter is computed using Monte Carlo integration as described above.

This estimator of γ_2 is also asymptotically normal, but its asymptotic variance is complicated by the estimation of $\hat{\boldsymbol{\gamma}}_1$ in the first step. Using the standard procedure to account for such nuisance parameters, it follows that the asymptotic variance of γ_2 may be estimated by²⁴

$$\text{Var}[\hat{\gamma}_2] = \frac{1}{N} \left(1 + \frac{1}{\pi_2}\right) (\hat{\mathbf{A}})^{-1} \text{Var}[f(\hat{\gamma}_2, u_i; \hat{\boldsymbol{\gamma}}_1)] (\hat{\mathbf{A}})^{-1} + (\hat{\mathbf{A}})^{-1} \hat{\mathbf{F}} \text{Var}[\hat{\boldsymbol{\gamma}}_1] \hat{\mathbf{F}}' (\hat{\mathbf{A}})^{-1} \quad (2)$$

where π_2 is a scaling factor controlling the number of simulations in the second estimation step. Here $\hat{\mathbf{A}} \equiv \frac{1}{\pi_2 N} \sum_{s=1}^{\pi_2 N} \frac{\partial m(\hat{\gamma}_2, u_s; \hat{\boldsymbol{\gamma}}_1)}{\partial \gamma_1}$, $\hat{\mathbf{F}} \equiv \frac{1}{\pi_2 N} \sum_{s=1}^{\pi_2 N} \frac{\partial m(\hat{\gamma}_2, u_s; \hat{\boldsymbol{\gamma}}_1)}{\partial \boldsymbol{\gamma}_1}$, and $\widehat{\text{Var}}[f(\hat{\gamma}_2, u_i; \hat{\boldsymbol{\gamma}}_1)]$ denotes the vari-

²³To avoid perfect linearity between the chosen moments and hence insure full rank of the Jacobian $\frac{\partial \mathbf{m}_1}{\partial \boldsymbol{\gamma}_1}$, we do not include the probability of giving more than 600 DKK $P(g_{it} > 600)$ in the matched moments.

²⁴The derivations are deferred to the Supplementary Appendix.

ance of $\frac{1}{N} \sum_{i=1}^N m_2(u_i)$ when re-centered around $E[m_2(\hat{\gamma}_2; \hat{\gamma}_1)]$. The first term of Equation (2) represents the standard variance from MSM, whereas the second term accounts for estimation uncertainty about $\hat{\gamma}_1$ from our first step.

6.4 Structural estimates

Table 5 shows that our model matches the low propensity to give in the empirical data and the empirical distribution of giving in different ranges quite well. However, we underestimate the probability of giving between 100 DKK and 200 DKK (0.0040 in the data vs. 0.0029 based on the model), and we overestimate the autocorrelation in the data (35.9 in the data vs. 40.9 implied by our model).

The unconditional mean of a_{it} is estimated to $\mu = -1,274$ and the standard deviation σ to 704. The low mean captures the fact that on average across the pre-treatment period 98.79% of individuals did not donate in any given month. Hence, a large share of the distribution of a is below the donation threshold \bar{a} . At the same time the distribution of a must capture the tendency that people most frequently give between 100 and 200 DKK conditional on donating. We note that the model implies that potential donors get zero donation utility if they do not give a strictly positive amount. Hence, a negative value of a does not translate into negative utility. We estimate the persistence in the process for a to 0.22, suggesting that the warm-glow parameter in the previous period has a modest but statistically significant positive effect on current period warm-glow.

The estimate of the annoyance cost Λ depends on the planning horizon T of the donors. When $T = 12$, donors have a one year horizon, and we estimate the annoyance cost to 12.95 DKK (\simeq \$2.35). This is quite similar in magnitude to the social pressure cost estimated in DellaVigna et al. (2012). The average hourly wage in Denmark was DKK 243 in 2012 (Danmarks Statistik, 2013) meaning that the annoyance cost equals about 5% of the average hourly wage.

Figure 6 illustrates how the estimated annoyance cost varies with the planning horizon T . The estimates increase sharply until $T \simeq 10$ after which it stabilizes at about 14 DKK. This suggests that there is little difference in the estimates for planning horizons beyond a year ($T = 12$).

Intuitively, the positive relationship between the planning horizon and the estimated annoyance cost can be explained as follows. If the potential donor has a very short planning horizon (i.e. T small), the option value of subscribing is small, because there are few future periods in which he could get a high realization of a_{it} and hence be generous enough to give. Hence, the potential donor

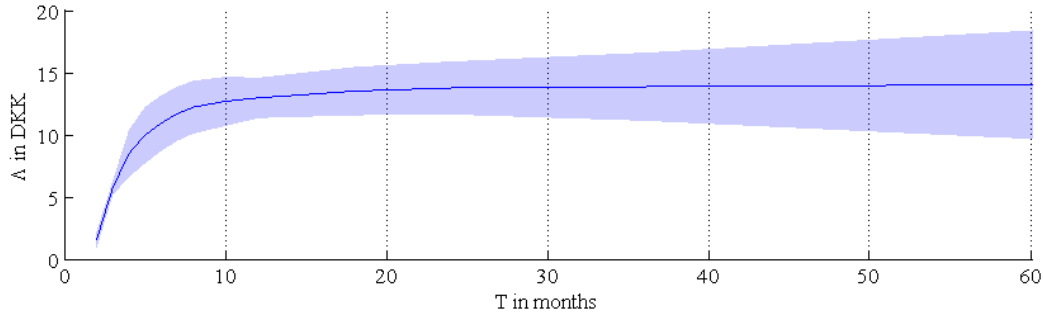
Table 5: Structural estimates

Estimated parameters		
Model parameters		
Mean weight on warm-glow, μ	-1,274.48 (19.211)	
Auto correlation in warm-glow, ρ	0.22 (0.001)	
Std.dev. of warm-glow error term, σ	704.49 (19.211)	
Annoyance cost, Λ , when $T = 12$	12.95 (0.805)	
	Model moments	Empirical moments
Step 1:		
$P(g_{it} = 0)$	0.9885	0.9879
$P(0 < g_{it} \leq 100)$	0.0036	0.0031
$P(100 < g_{it} \leq 200)$	0.0026	0.0040
$P(200 < g_{it} \leq 400)$	0.0030	0.0029
$P(400 < g_{it} \leq 600)$	0.0014	0.0012
$E(g_{it}g_{it-1})$	40.930	35.906
N	249,380 ^a	12,469
\mathcal{T}	54	54
Step 2:		
$P(u_i = 1)$	0.0050	0.0050
N	3,000,000 ^a	12,469

Notes: The upper part of the table reports estimated parameters with standard errors in brackets. The bottom part of the table reports model-implied and empirical moments. Empirical moments are calculated for individuals in Control II for whom data on donation history is available. Model-implied moments are calculated from simulated data. a) reports the total number of simulations, i.e., $N\pi$.

will be relatively more likely to unsubscribe, and a relatively low annoyance cost Λ is needed to explain the unsubscription rate in our data. On the other hand, if the potential donor has a long planning horizon, the option value of subscribing is larger because there are more periods in which the individual can potentially get a high enough realization of a_{it} to make a donation. In this case a relatively high value of Λ is needed to explain the empirical unsubscription rate.

Figure 6: Estimated annoyance cost Λ as a function of the planning horizon T



Notes: The solid line shows estimates of the annoyance cost Λ in DKK for different planning horizons T . The shaded area illustrates the 95% confidence interval for the point estimates.

7 Welfare implications of reminders

This section conducts a welfare analysis to quantify the relative size of the costs and benefits of reminders. We begin by estimating the effect of unsubscribing in Experiment I on giving. These results and the structural estimates of the model provided in Section 6 are then used in our welfare analysis.

7.1 Effect of unsubscribing on giving behavior

The first question we want to address is whether the charity loses valuable donors through unsubscriptions. If only marginal donors or individuals who donate just as much without a reminder unsubscribe, then a high frequency of reminders is not problematic and might even be desirable for the charity. For the sample in Experiment I we have access to donation data including amount and time of giving for a reasonably long period before and after the treatment period. This makes a difference-in-difference analysis of donation behavior for subscribers and unsubscribers possible.

Since unsubscribing only means that donors switch off *one* communication channel, it is possible that unsubscribers will donate even after they have unsubscribed. As illustrated by the time series in Figure 7, unsubscribers appear to be relatively more marginal donors in terms of giving than subscribers both before and after the treatment month. The difference-in-difference analysis in Table 6 confirms this result as unsubscribers are less likely to give, and the unconditional average amount given per month is smaller for unsubscribers than subscribers in both the pre-

and post-treatment period. However, the differences in the amounts donated are not statistically significant.

The likelihood of donating decreases for both subscribers and unsubscribers from the pre-treatment period to the post-treatment period, and this results appears to be largely driven by a drop in donations around Christmas (Figure 7). The decrease in Christmas giving appears larger for the unsubscribers than for the subscribers, and overall we find a larger drop in the propensity to give and the unconditional amount given for unsubscribers than for subscribers. The difference-in-difference estimate of the effect of unsubscribing on the monthly amount donated is 0.62 DKK. For the propensity to give the corresponding estimate is 0.003 (Table 6). However, these differences are not statistically significant.²⁵

We note that there may be a couple of reasons why the difference-in-difference estimates are insignificant. First, a relatively small share of our sample are unsubscribers, and this reduces the power of our test. Second, the e-mails in the first experiment were among the first e-mails that individuals in our sample got from the charity and hence represent one of the first opportunities to unsubscribe. The treatment period may therefore only work as a rather “fuzzy” treatment period in terms of unsubscriptions because some unsubscribers may have mentally unsubscribed before the treatment period.

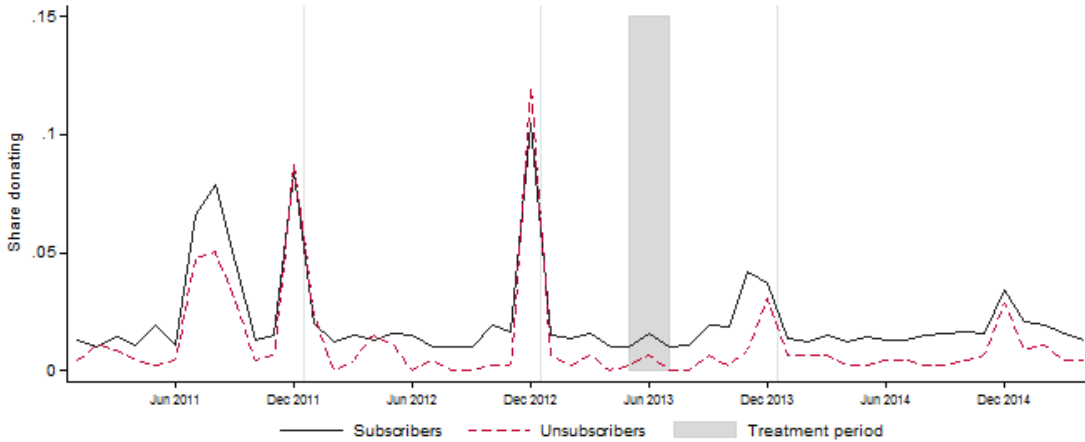
7.2 Welfare of potential donors and the charity

We next evaluate the welfare consequences of sending out a regular e-mail to remind people to donate, i.e., the welfare effects of the e-mail sent to the control group in Experiment II.²⁶ Potential donors either give or do not give depending on their realization of a_{it} when they receive the reminder. Donors with $a_{it} < 1$ do not give and have a negative utility because they incur the annoyance cost Λ without the warm-glow from giving. Donors with $a_{it} > 1$ give and receive warm-glow utility, but they also incur the annoyance cost. Overall potential donors might incur positive or negative utility from receiving the e-mail. Conditional on giving in response to the message, the welfare of the donor is 723 DKK on average (see Table 7). For individuals who do not give, the

²⁵Unsubscribers in the the Targeted Reminder and Control I treatments are not different in their likelihood of giving in any period (results available on request). Hence, we do not have evidence to suggest that unsubscribers in the Targeted Reminder treatment are more or less marginal on average than unsubscribers to the first e-mail.

²⁶We consider the effect of reminders in Control II because the communication this group received was closest to the usual communication sent by the charity.

Figure 7: Share donating per month by unsubscription behavior in Experiment I



Note: Based on Experiment I sample and including only individuals for whom we have donation history (14,995 subscribers and 458 unsubscribers, i.e., 93.1% of the sample). Data covers the period 1 Jan 2011 - 31 Mar 2015.

Table 6: Analysis of propensity to give and unconditional amount donated

	Subscribers (1)	Unsubscribers (2)	DID ^a (1)-(2)	DID with controls ^{a,b}
Propensity to give per month				
Pre-period	0.024 (0.086)	0.016 (0.021)	-0.008* (0.004)	
Post-period	0.018 (0.098)	0.007 (0.042)	-0.011** 0.004	
Change in propensity to give	-0.006*** (0.001)	-0.009*** (0.002)	-0.003 (0.006)	-0.006 (0.011)
Unconditional amount given DKK/month				
Pre-period	7.78 (0.24)	6.71 (0.93)	-1.07 (1.44)	
Post-period	5.15 (0.27)	3.46 (1.80)	-1.69 (1.44)	
Change in mean amount given	-2.63*** (0.35)	-3.25 (2.02)	-0.62 (2.04)	-1.39 (3.82)

Notes: The table reports means and standard errors (in brackets). Includes Experiment I sample only and only the 93.1% percent of the sample we have donation history for. Data covers the periods 1 Jan 2011 - 31 Mar 2015. The treatment period was Jun 2013. ^a means and standard errors are estimated by linear regression with data collapsed into one pre-treatment and one post-treatment period. We use one pre- and one post-treatment period to address concerns about reliable estimation of standard errors as discussed by Bertrand et al. (2004). ^b the regression includes the following controls: age, female dummy, number of months since last donation, and amount donated last time. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

welfare is -12.95 DKK. On average 1.53% of donors give in a given month. Hence, in expectation the message diminishes the welfare of potential donors by 1.69 DKK.

For the charity there is a positive immediate effect of sending the e-mail which is equivalent to the additional amount raised, but there is also a long-term cost of lost future donations from unsubscribers. We use the difference-in-difference results in Table 6 to quantify the potential long-term effects and estimate that the charity loses 0.62 DKK per unsubsubscriber per month. The net effect for the charity depends on how strongly the charity discounts future periods and for how many periods donations are lost. In our setting, unsubscribing is an absorbing state, i.e., once people have unsubscribed they cannot rejoin the list with the same e-mail address. We therefore assume that the effect lasts for the remaining lifetime of the potential donors. Since the negative effects of unsubscribing for the charity occurs over future periods, the net effect depends on the charity's discount rate. We show results for discount rates of 2%, 10%, and 20%. The net effect for the charity is generally small but positive - for instance 1 DKK for a discount rate of 10%.

Our results suggest that the negative welfare effect for potential donors is greater than the net effect for the charity. This holds even if the negative long-term effects for the charity are set to zero. Hence, the money raised by the charity must be used very efficiently to generate positive welfare effects.²⁷

Table 7: Welfare effect of a regular fund-raising e-mail

Potential donors			
Welfare conditional on giving			722.69
Percentage of potential donors who give (%)			1.53
Welfare conditional on not giving			-12.95
<i>Expected welfare effect per contacted individual</i>			<i>-1.69</i>
Charity			
Charity discount rate (%)	2%	10%	20%
Immediate effect: Money raised per potential donor	1.38	1.38	1.38
Long-term effect: Money lost per potential donor	1.13	0.38	0.20
<i>Net fund-raiser effect per contacted individual</i>	<i>0.25</i>	<i>1.00</i>	<i>1.18</i>

Notes: Figures are in DKK unless otherwise stated. Long-term effects for the charity are assumed to last for the remaining lifetime of the potential donors. On average donors in the sample are 38 years old and have a life expectancy of 86 years.

²⁷This could be achieved if we make the assumption that a dollar given to a poor person in a developing country or invested in necessary infrastructure creates larger welfare effects than the welfare the same dollar could have created in the pocket of a potential donor (Singer, 2009).

An important caveat to our analysis is that our estimates do not account for wider effects which may be either positive or negative. For example, there may be a substitution effect implying that unsubscribers do not donate less in total after unsubscribing but instead substitute donations to another charity. This would bias upwards the welfare calculations. However, complementary effects might also arise where potential donors start paying less attention to the content of messages and reminders when they receive more of them. This could have negative welfare consequences as it reduces the positive effect of reminders and increases the annoyance costs.

One obvious remedy reducing annoyance costs is a better allocation of reminders. Rather than sending e-mails to every donor, the charity could provide donors with an opportunity to self-select into the reminder. This could however lead to very low sign-up rates due to inertia. So a default enrollment is preferred (Thaler and Sunstein, 2003). As is common in fundraising, our sample of individuals had consented to the charity storing their e-mail address when they made a donation. A better solution than an opt-in could be to allow individuals to adjust the frequency of the newsletter. As donations spike around Christmas and New Year, some individuals might only want to be reminded around Christmas to give. Alternatively, the charity could use the available giving data of the donors to predict when certain individuals are most likely to give and send targeted reminders at those points.²⁸ The availability of more detailed data through an increase in online giving could make this a feasible strategy for charities and organizations in the near future. On the other hand, natural catastrophes or crises do not follow a predictable pattern. A large crisis such as a hurricane could create an external shock to a person's a_{it} , which would make him willing to make a donation at another time of the year than what would be predicted by past donation behavior. Nevertheless, a more precisely targeted approach to fundraising would help balance the social welfare costs.²⁹

²⁸Allcott and Kessler (2015) have a similar discussion for the opt-in into Home Energy Reports (HERs) and for sending targeted reports based on the customers willingness to pay for HERs.

²⁹The literature on optimal catalog mailings (among others Simester et al. (2006); Gönül and Shi (1998); Gönül and Hofstede (2006)) has also advocated that companies should take long run implications into account. Papers in this tradition argue that companies might be able to increase profits by incurring the immediate cost of mailing a catalog to recipients who are not expected to make a purchase in the short term because some recipients can be expected to purchase in the longer term. However, our argument is that one might then overlook the long run annoyance costs to the recipient which leads to loss of consumer base for the advertiser.

8 Conclusion

This paper documents the hidden costs of nudging in the context of reminders to subscribers on a charity’s mailing list. Our results show that reminders increase the number of donations and at the same time increase the number of unsubscriptions from the mailing list. We explore the reasons for unsubscribing in a theoretical model which we then test in a second field experiment. Based on data from the second experiment, we also structurally estimate the annoyance costs and the utility of giving. Finally, we conduct a welfare analysis from the perspective of the subscribers and for the charity. We find that the annoyance cost of a reminder is about 13 DKK each month, and on average a reminder is welfare diminishing for potential donors. Furthermore, when accounting for the long-term effects of unsubscriptions on giving, the net effect for the charity of sending a reminder is just 1 DKK.

The model we develop and test experimentally in this paper could be extended to other settings where reminders are used to tackle inattention or procrastination. Instead of warm-glow from donating, the benefits could be improved health, savings, or academic outcomes. It is easy to see that the higher the personal benefit of the reminder and the smaller the cost of the prompted action, the larger the utility from the reminder, irrespective of the potential annoyance costs. However, high frequency or very pushy reminders create a welfare diminishing cost even in these settings.

Concluding, we recommend that the literature on nudging policies continues to include more critical evaluations of policies by pricing in the psychological costs of the nudge and unintended behavioral reactions and include them in the welfare calculations for all affected parties.

Appendices

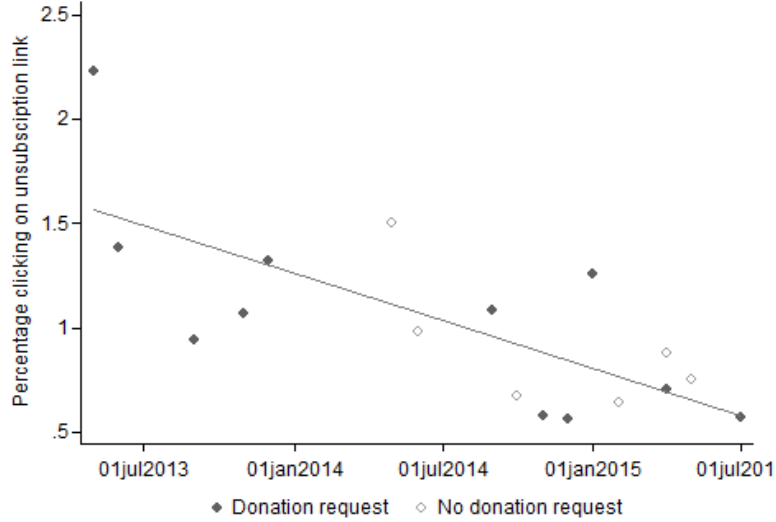
A Figures and tables

Table A1: Amount donated unconditionally in Experiment I and II

	Experiment I		Experiment II	
	(1)	(2)	(3)	(4)
Targeted Reminder	0.42196 (0.31184)	0.61872 (0.69962)		
Low Frequency			-0.24689 (0.25171)	-0.37584 (0.28892)
Future Benefit			-0.26455 (0.25335)	-0.31816 (0.30103)
Female		0.38122 (0.88231)		0.39175* (0.23077)
Age		0.04882* (0.02570)		0.08134*** (0.01078)
City		-0.01183 (0.59964)		0.70983*** (0.24663)
Months since last donated		-0.04957*** (0.01594)		-0.01337* (0.00753)
Amount last donated		0.00099 (0.00105)		0.00098** (0.00041)
Months on e-mail list				-0.09101*** (0.03329)
Constant	0.62701*** (0.20449)	-0.09647 (1.29571)	1.37191*** (0.21369)	0.47054 (0.82847)
Observations	17391	6448	43489	27053
R^2	0.00	0.00	0.00	0.01

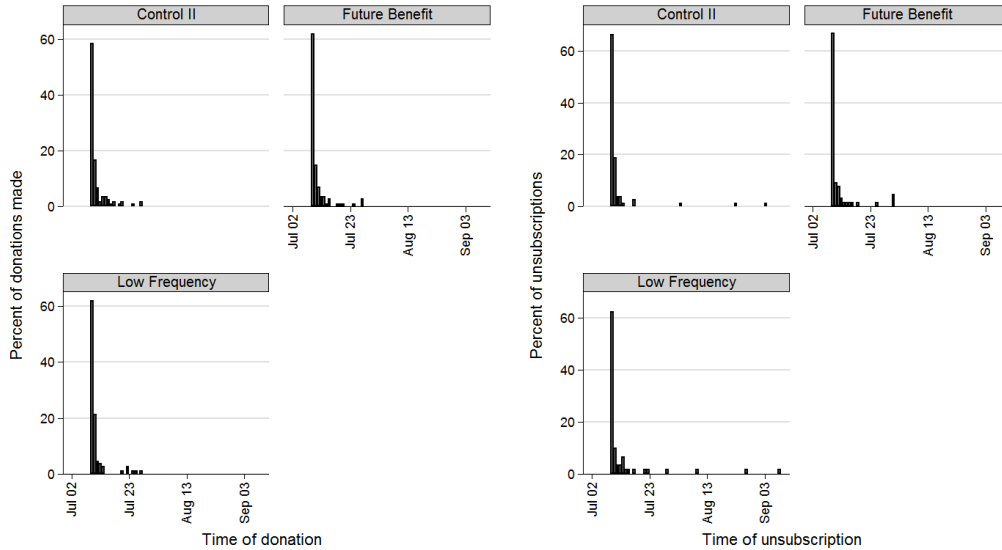
Notes: The table shows OLS regressions on the amount donated unconditional on donating with standard errors in parentheses. Targeted Reminder, Low Frequency, and Future Benefit are dummy variables that are evaluated in comparison to their respective control groups. Female and City are dummy variables. Months since last donated and Amount last donated correspond to the last donation prior to the respective experiment through any channel. Months on e-mail list is set at one month for everyone in the first experiment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure A1: Unsubscription rates over time



Notes: The percentage of potential donors who clicked on the unsubscription link in a random selection of previous e-mails.

Figure A2: Timing of giving and unsubscriptions in Experiment II

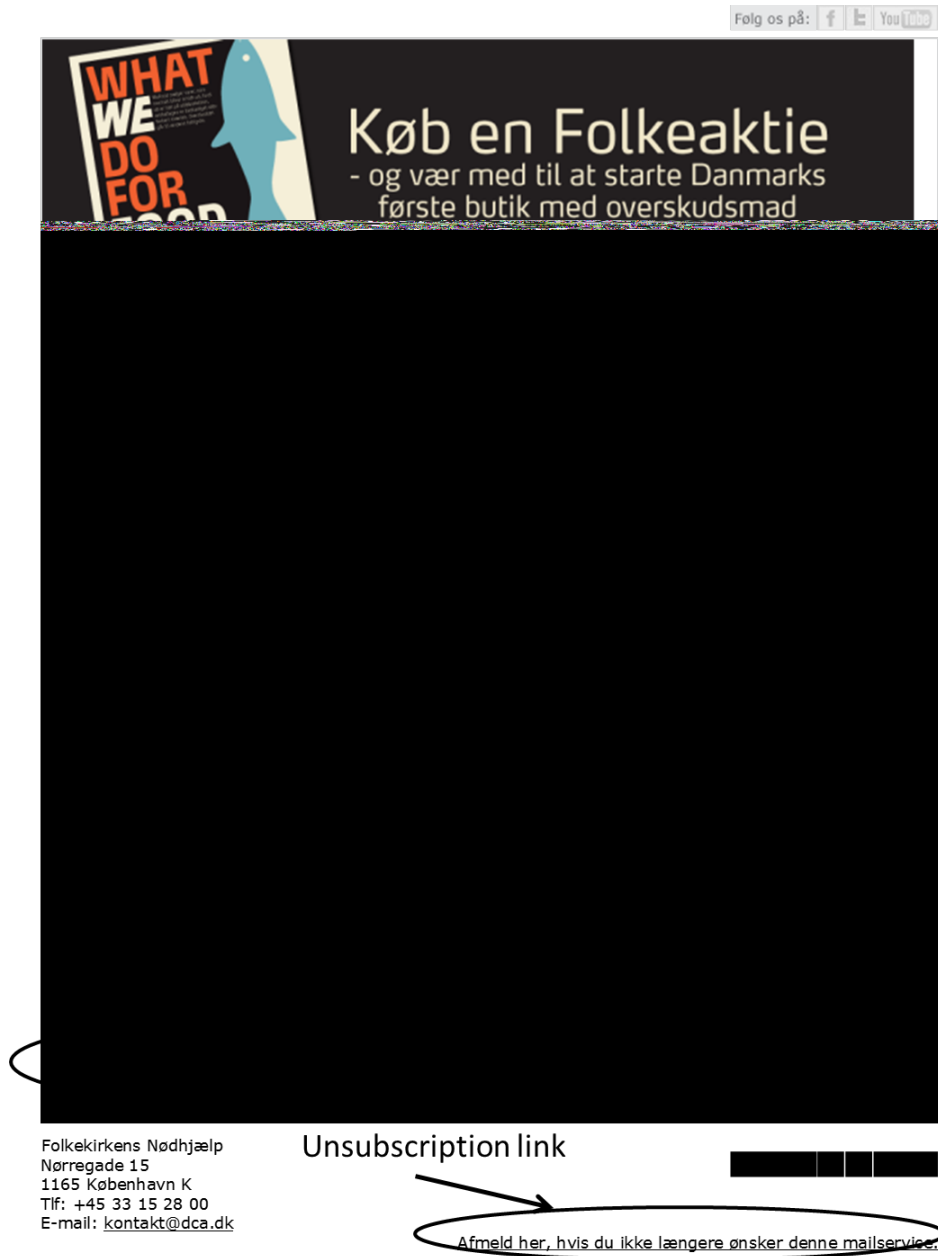


Notes: The figure shows the percentage of donations and unsubscriptions made on each day in the Control II, the Low Frequency, and the Future Benefit treatments. The e-mail was sent on July 9th for all three groups.

Figure A3: Experiment I: First fundraising e-mail



Figure A4: Experiment II: Control July 2015 e-mail



B Reasons for unsubscribing

From the landing page for unsubscribers in Experiment II, we gather some indicative information on why individuals unsubscribe from the mailing list. Figure B1 provides information on reasons for unsubscribing and shows answers given on the unsubscription page for the three treatments combined.³⁰ The most common reason for unsubscribing is that recipients do not want to give to the charity anymore, followed by a dislike of being asked directly for a donation and “Other” reasons.³¹ The second most frequently chosen reason is that DCA sends out too many e-mails. Clearly we observe self-selection into answering the questions as a result of unsubscribing. Less people unsubscribed in the Low Frequency treatment, which means that we observe a smaller group of unsubscribers, who answer this question. Therefore, we cannot make any causal inference from the data, but just use it as a glimpse into why individuals unsubscribe from the mailing list. In section C we compare the reasons given in Experiment II to those given in the next solicitation e-mail.

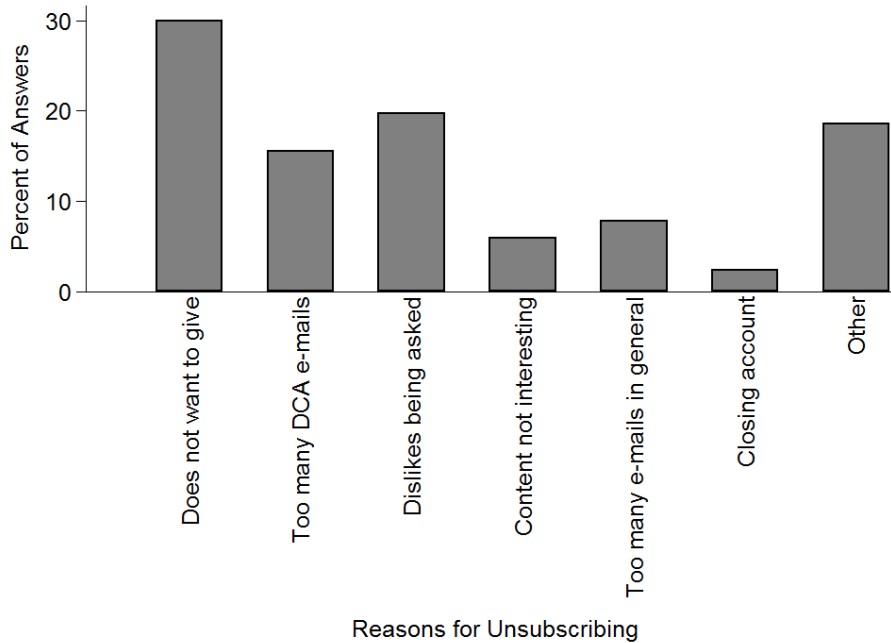
C Medium term effects of Experiment II

Four weeks after we ran Experiment II, on the 14th of August 2015, we sent out the announced matching opportunity to all individuals who had been part of Experiment II, except those who had unsubscribed. This mailing had been announced to the individuals in the Future Benefit treatment, but to the other two treatment groups it came as a surprise. This e-mail took the upcoming school start in Denmark as a reason to ask subscribers for a donation of a school bench worth 200 DKK to some of the world’s poorest children. For one week, until the 20th of August, we added a matching gift consisting of a school meal (worth 20 DKK) to a poor child for everyone that donated 200 DKK during that period. The second e-mail had the subject line “School start”. Due to the unsubscriptions from all three groups and the expectations formed by the Future Benefit group about the content of this e-mail, we have possible self-selection into opening the new mailing. The following results thus have to be interpreted with the selection effect in mind. They nevertheless give an insight into medium term effects of the experimental treatments on subsequent behavior.

³⁰We pooled the answers, as there are no significant differences between the treatment regarding the reasons for unsubscribing.

³¹Since it was technically impossible to randomize the answer possibilities, the first reason might be slightly overstated due to order effects.

Figure B1: Reasons for unsubscribing in Experiment II



Notes: The figure shows the percentage of each reason given for unsubscribing when combining the three treatments.

While the response rate in terms of number of donations was lower than in Experiment II, the average donation conditionally on donating was higher. Forty-eight individuals donated within three days in response to the mailing (see Table C1). That is 0.14 percent of the total group. Conditionally on donating the average amount donated was 258 DKK which is higher than in Experiment II. The differences are likely due to the solicitations being for different public goods (a food waste shop in Denmark vs. school benches for poor children) and due to a higher focal amount (in Experiment II this was 100 DKK, and in the subsequent mailing it was 220 DKK). The unsubscription rate was 0.66 percent and hence slightly higher than in Experiment II.

Mostly, we are interested in whether the treatments in Experiment II had any effects on the response to the mailing four weeks later. Out of the 259 receivers that unsubscribed after receiving the mailing, 70 were in the Control II group, 97 in the Future Benefit group, and 90 in the Low Frequency group. Chi2 tests on the differences reveal that significantly more people from the Future Benefit treatment than the Control II group unsubscribed in response to the subsequent

Table C1: Follow-up mailing: Results statistics

	All	Control II	Low Frequency	Future Benefit
Responses in (%)				
Percentage who gave	0.12 (3.41)	0.12 (3.52)	0.10 (3.21)	0.10 (3.21)
Percentage who unsubscribed	0.60 (7.66)	0.48 (6.92)	0.62 (7.84)	0.67 (8.14)
Observations				
Full sample	43,292	14,450	14,428	14,414
Number who gave	48	18	15	15
Number of unsubscribers	259	70	90	97

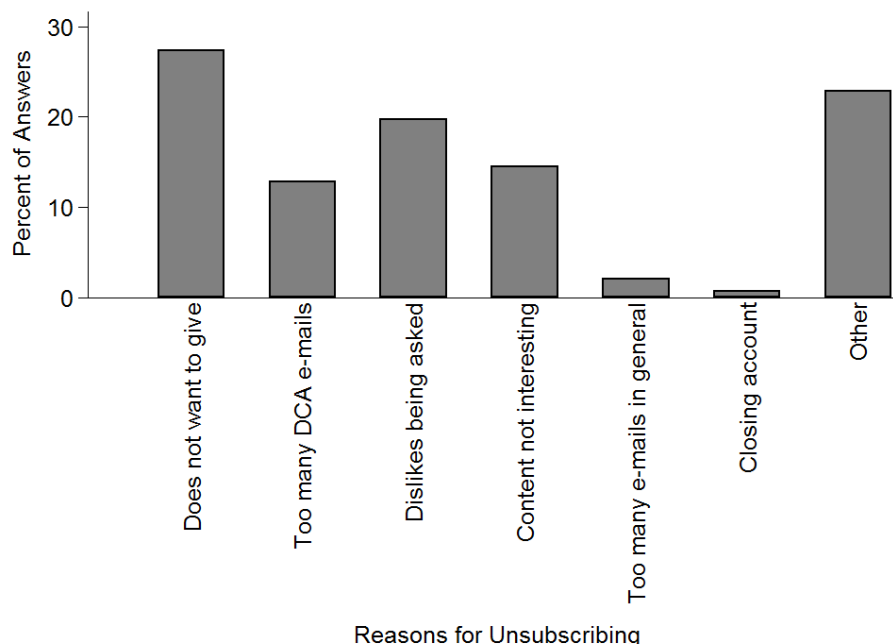
Notes: The table provides means and standard deviations (in brackets). The response rates are given in percent. The table also provides information about the total number of observations in each sample and the total number of donations.

mailing (Chi2, p-value=0.04). The other differences in unsubscription rates are not significant.³² This result could be due to a higher self-selection into reading the e-mail for individuals in the Future Benefit group, as the matching opportunity was announced in their e-mail in Experiment II. A second explanation could be a disappointment because the offered match and the cause does not meet expectations created in the treatment e-mail. In order to understand which explanation is more likely, we look at the stated reasons for unsubscribing.

We do not see a difference in the distribution of reasons for unsubscribing between the Future Benefit and the other two treatments. This lends credibility to the explanation that rather than being disappointed more individuals in the Future Benefit group took the time to open the e-mail and formally unsubscribe instead of just deleting it without opening it. In the “Other” comments we do not observe any reference to the matching possibility or any form of potential disappointment with the offer. We display the pooled distribution in Figure C1. Overall, it is interesting to note that again “Dislike being asked” was selected more often than “Too many DCA e-mails”. This difference actually seems understated as among the “Other” responses several were along the lines “I decide myself, when I want to give” or “I give when I can, but not when you ask”. Thus, our results from both solicitations support the “avoiding the ask” theory (Dana et al., 2007; Andreoni

³²The p-value for a Chi2 test of differences in the unsubscription rate for the two treatments is 0.61, and the p-value for a test for difference in unsubscription rates in the Low Frequency and Control II group is 0.11.

Figure C1: Reasons for unsubscribing in first solicitation after Experiment II



Notes: The figure shows the percentage of each reason given for unsubscribing when combining the three treatments.

et al., 2011; DellaVigna et al., 2012; Knutsson et al., 2013), but should be interpreted with caution because of the self-selection.

We do not observe a higher rate of giving for the Future Benefit group, which could have been expected as a result of more people opening the e-mail. Eighteen individuals of the control group, 15 of the Future Benefit, and 15 of the Low Frequency treatment donated in this mailing. Four individuals donated both in Experiment II and in the subsequent campaign, and they were divided across all the treatments (1/2/1). Since the treatments in Experiment II had no effect on the number on donations in the experiment, we do not expect a substitution effect in the subsequent mailing. From the response rate, it is obvious without formal testing that announcing the match in the treatment e-mail did not have a significant positive effect on subsequent giving. Because of the already mentioned self-selection issue, we will not go further in our interpretation of the results.

References

- Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics* 156, 38–67.
- Allcott, H. and J. B. Kessler (2015, October). The welfare effects of nudges: A case study of energy use social comparisons. Working Paper 21671, National Bureau of Economic Research.
- Allcott, H. and T. Rogers (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review* 104(10), 3003–37.
- Altmann, S. and C. Traxler (2014). Nudges at the dentist. *European Economic Review* 72(0), 19 – 38.
- Andreoni, J. (1989). Giving with impure altruism: Applications to charity and ricardian equivalence. *Journal of Political Economy* 97(6), 1447–1458.
- Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *Economic Journal* 199(401), 464–477.
- Andreoni, J., J. M. Rao, and H. Trachtman (2011). Avoiding the ask: A field experiment on altruism, empathy, and charitable giving.
- Bernheim, B. D., A. Fradkin, and I. Popov (2015). The welfare economics of default options in 401(k) plans. *American Economic Review* 105(9), 2798–2837.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics* 119(1), 249–275.
- Bhargava, S. and D. Manoli (2015). Psychological frictions and the incomplete take-up of social benefits: Evidence from an irs field experiment. *American Economic Review* 105(11), 3489–3529.
- Bhattacharya, J., A. M. Garber, and J. D. Goldhaber-Fiebert (2015). Nudges in exercise commitment contracts : a randomized trial. Parallel als Online-Ausg. erschienen.

- Cain, D. M., J. Dana, and G. E. Newman (2014). Giving versus giving in. *The Academy of Management Annals* 8(1), 505–533.
- Calzolari, G. and M. Nardotto (2014). Nudging with information: A randomized field experiment. Working paper.
- Card, D., S. DellaVigna, and U. Malmendier (2011). The role of theory in field experiments. *Journal of Economic Perspectives* 25(3), 39–62.
- Carr, P., R. Jarrow, and R. Myneni (1992). Alternative characterizations of american put. *Mathematical Finance* 2(2), 87–106.
- Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009). Optimal defaults and active decisions. *The Quarterly Journal of Economics* 124(4), 1639–1674.
- Chesterley, N. (2015). Nudges and welfare: Optimal behavioural policy with decision costs. *SSRN Working Paper*.
- Conlin, M., T. O’Donoghue, and T. J. Vogelsang (2007). Projection bias in catalog orders. *American Economic Review* 97(4), 1217–1249.
- Damgaard, M. T. and C. Gravert (2014). Now or never! the effect of deadlines on charitable giving. *Aarhus University Economics Working Paper 2014-03*.
- Dana, J., R. A. Weber, and J. X. Kuang (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* 33(1), 67–80.
- DanChurch Aid (2013). Aarsrapport folkekirkens nødhjælp 2013.
- Danmarks Statistik (2013). Befolkningens løn.
- DellaVigna, S., J. A. List, and U. Malmendier (2012). Testing for altruism and social pressure in charitable giving. *Quarterly Journal of Economics* 127(1), 1–56.
- Eckstein, Z. and K. I. Wolpin (1989). The specification and estimation of dynamic stochastic discrete choice models: A survey. *Journal of Human Resources* 24(4), 562–598.

- Gilbert, B. and J. G. Zivin (2014). Dynamic salience with intermittent billing: Evidence from smart electricity meters. *Journal of Economic Behavior and Organization* 107, Part A(0), 176 – 190.
- Gönül, F. and F. T. Hofstede (2006). How to compute optimal catalog mailing decisions. *Marketing Science* 25(1), 65–74.
- Gönül, F. and M. Z. Shi (1998, September). Optimal mailing of ccatalog: A new methodology using estimable structural dynamic programming models. *Management Science* 44(9), 1249–1262.
- Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review* 103(7), 2643–82.
- Harrison, G. W. (2014). Cautionary notes on the use of field experiments to address policy issues. *Oxford Review of Economic Policy* 30(4), 753–763.
- Huck, S. and I. Rasul (2010). Transaction costs in charitable giving. *The B.E. Journal of Economic Analysis & Policy* 10(1), Article 31.
- ISOBRO and Deloitte (2014). Isobro og deloitte's analyse af udviklingen i de danske indsamlingssorganisationers indtægter mv. i perioden 2009 - 2013.
- Jacka, S. D. (1991). Optimal stopping and the american put. *Mathematical Finance* 1(2), 1–14.
- Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman (2012). Getting to the top of mind: How reminders increase saving. *Working paper*.
- Kim, I. J. (1990). The analytical valuation of american options. *Review of Financial Studies* 3(4), 547–572.
- Knutsson, M., P. Martinsson, and C. Wollbrant (2013). Do people avoid opportunities to donate?: A natural field experiment on recycling and charitable giving. *Journal of Economic Behavior and Organization* 93, 71 – 77.
- Laibson, D., A. Repetto, and J. Tobacman (2007). Estimating discount functions with consumption choices over the lifecycle.

- Macharia, W. M., G. Leon, B. H. Rowe, B. J. Stephenson, and B. Haynes (1992). An overview of interventions to improve compliance with appointment keeping for medical services. *Journal of the American Medical Association* 267(13), 1813–1817.
- McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica* 57(5), 995–1026.
- Murooka, T. and M. A. Schwarz (2016). The timing of choice-enhancing policies. *Unpublished Manuscript*.
- Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica* 55(5), 999–1033.
- Rust, J. (1994). *Handbook of Econometrics*, Volume 4, Chapter Structural Estimation of Markov Decision Processes, pp. 3081–3142. North Holland, Amsterdam: Elsevier Science B.V.
- Simester, D. I., P. Sun, and J. N. Tsitsiklis (2006, May). Dynamic catalog mailing policies. *Management Science* 52(5), 683–696.
- Singer, P. (2009). *The life you can save: Acting now to end world poverty*. Random House.
- Sonntag, A. and D. J. Zizzo (2015, 08). On reminder effects, drop-outs and dominance: Evidence from an online experiment on charitable giving. *PLoS ONE* 10(8), e0134705.
- Taubinsky, D. (2013). From intentions to actions: A model and experimental evidence of inattentive choice. *working paper*.
- Thaler, R. H. and C. R. Sunstein (2003). Libertarian paternalism. *American Economic Review* 93(2), 175–179. Papers and Proceedings.
- Trachtman, H., A. Steinkruger, M. Wood, A. Wooster, J. Andreoni, J. Murphy, and J. Rao (2015). Fair weather avoidance: unpacking the costs and benefits of avoiding the ask? *Journal of the Economic Science Association* 1(1), 8–14.
- Vervloet, M., A. J. Linn, J. C. M. van Weert, D. H. de Bakker, M. L. Bouvy, and L. van Dijk (2012). The effectiveness of interventions using electronic reminders to improve adherence to chronic medication: A systematic review of the literature. *Journal of the American Medical Association*, 19, 696–704.