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

Mette Trier Damgaard† Christina Gravert‡
Aarhus University University of Gothenburg

September 2, 2017

Abstract

We document the hidden costs of a popular nudge and show how these costs distort policy making when neglected. In a field experiment with a charity, we find reminders increasing intended behavior (donations), but also increasing avoidance behavior (unsubscriptions from the mailing list). We develop a dynamic model of donation and unsubscription behavior with limited attention. We test the model in a second field experiment which also provides evidence that the hidden costs are anticipated. The model is estimated structurally to perform a welfare analysis. Not accounting for hidden costs overstates the welfare effects for donors by factor ten and hides potential negative welfare effects of the charity.

Keywords: Avoiding-the-ask, charitable giving, field experiment, inattention, nudge, reminders.
JEL codes: C93, D03, D64, H41

*We are grateful to the editor and two anonymous referees as well as Stefano DellaVigna, DanChurchAid, Martin M. Andreasen, Ned Augenblick, Teodora Boneva, Alexander Koch, John List, Ulrike Malmendier, Takeshi Murooka, Georg Nöldeke, Heiner Schuhmacher, Robert Sudgen and several seminar and conference participants for comments and suggestions. Funding by Swedish Research Foundation, Købmand Ferdinand Sallings Mindefond and Aarhus University. AEA RCT Registry 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, Vasagatan 1, 41124 Gothenburg, Sweden, e-mail: christina.gravert@economics.gu.se

1 Introduction

"Nudging" policies have gained increased attention from practitioners and academics. Nudges are small deliberate changes to the decision environment designed to increase privately and socially beneficial behavior such as healthy habits, increased saving, sustainable consumption or charitable giving without adjusting prices or restricting choice. With the establishment of governmental behavioral units like in the UK, the US and Denmark behavioral interventions are becoming part of the policy toolkit.

The success of a nudge is usually evaluated by the positive behavior change it induces. Moreover, their low implementation costs create a high cost-benefit ratio. 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 (see Carroll et al. (2009); Handel (2013); Allcott and Kessler (2015); Bernheim et al. (2015); Bhattacharya et al. (2015); Chesterley (2015); Murooka and Schwarz (2016) for related arguments).

This paper investigates the cumulative welfare effects of nudges and shows that nudges might be less innocuous than generally assumed. We apply our approach to a well-known nudge: reminders.

Reminders are designed to curb forgetfulness by bringing a particular decision or task to recipients' attention and 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, 2016), adherence to medical treatments (Vervloet et al., 2012; Altmann and Traxler, 2014), personal savings (Karlan et al., 2016), 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 implementation costs of reminders, implying that reminders will become even more common. This makes it relevant to explore potential indirect or non-pecuniary costs, i.e. "hidden cost of nudging", of reminders for the recipients or the senders. We examine the aggregate reminder effect in the context of charitable giving. The hidden costs are identified from a revealed preference measure: 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 (Dana et al., 2007; Andreoni et al., 2017; DellaVigna et al., 2012; Knutsson et al., 2013; Cain et al., 2014; Trachtman et al., 2015). Every period, individuals decide whether to give or not if reminded about the donation possibility (Andreoni, 1989, 1990). 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 information and opportunities to donate. 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. 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 by two thirds, but also significantly increases unsubscriptions from the mailing list by a similar extent.

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 approximately 43,000 previous donors receives a regular solicitation e-mail from the charity. 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, instead of the regular monthly one. 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 who donates in response to the next e-mail. 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 by 39%. Announcing a future matching opportunity also reduces the number of unsubscriptions, but this result is only marginally significant. The treatments have no effect on the decision to donate or the donated amount, consistent with our model.

Using data from the second experiment, we structurally estimate the annoyance costs of being solicited through e-mail. Adding a structural estimation to the reduced form results enables us to study the underlying annoyance costs and provides inputs for a welfare analysis. We identify parameters influencing giving behavior from panel data on donations by individuals in the second experiment and we pin down the annoyance cost from the unsubscription rate. For the potential donor, there is a trade-off between annoyance costs and warm-glow from giving. We estimate the costs associated with *receiving* a reminder to 12.95 DKK ($1.95). This negative amount is on average slightly outweighed by the benefits of warm-glow from donating leading to an average welfare gain of 1.50 DKK ($0.23). However, by failing to consider the annoyance costs, a standard welfare evaluation would overstate the benefits of the reminder by a factor of ten.

We then consider the perspective of the charity by estimating the impact of a reminder on donations. 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 1.33 DKK ($0.20) per potential donor when using a discount rate of 10%. With a discount rate of 2% the net effect for the charity is negative.

The findings have important implications for public policy. The increasing volume of reminders, fueled by the encouraging results of previous studies, creates heretofore unanticipated costs for both receivers and senders. A one sided and short-term analysis based solely on the intended behavioral outcome, as is common today, can lead to negative surprises in the long-run. We encourage academics and policy makers to pay more attention to overall welfare effects.

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]

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 $t \in \{1,2,...,T\}$, 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 $v(g_t)$.[2] We model the cost of giving by the function $c(g_t)$ 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

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

where a_t is the weight on warm-glow utility.[3] The law of motion for a_t is given by an AR(1) process

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

where $\varepsilon_t \sim IID(0, \sigma^2)$ on a finite support $[-M, M]$, i.e., $M < \infty$. The AR(1) process introduces time-variation in the weight on warm-glow utility which can capture both variations in warm-glow from different fundraising campaigns and variations in the cost of giving, e.g. due to time-varying opportunity costs.[4]

[1]The technical details, including proofs, are provided in Online Appendix.

[2]Note 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.

[3]We assume that the warm-glow and cost functions are well-behaved i.e. $v'(\cdot) > 0$, $v''(\cdot) < 0$, $\lim_{g_t \to \infty} v'(g_t) = 0$ and $c'(\cdot) > 0$. 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.

[4]We also note that a deterministic process for a_t would lead to a static problem where the donor

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 of Karlan et al. (2016) and Taubinsky (2013), we assume that the donor is sophisticated and therefore aware of his inattention.[5]

We assume that *any* message from the charity serves as a reminder of the donation problem because it brings the charity back to the mind of the donor. The message does not need to refer back to a previous message, but can stand individually. It is also not necessary that the message contains an explicit ask. 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 with probability θ.[6] Hence, subscribing to the mailing list at the beginning of period t, increases the probability that the donor remembers the donation problem. If the donor is not subscribed he will miss out on future warm glow with probability $1-\theta$.

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 to a first approximation a moral cost of feeling guilty for having to be reminded.[7] We refer to this cost as an "annoyance cost", which for simplicity is assumed to be constant across people and across time to make the model tractable. We also assume that any type of message generates the same fixed cost, i.e., original solicitations and

either unsubscribes in period $t = 1$ or never unsubscribes.

[5]However, we note that the model does allow for overconfidence about prospective memory for donation decisions, as studied by for example Ericson (2011) and Letzler and Tasoff (2014), by letting θ being interpreted as the donor's subjective belief about the likelihood of remembering if not reminded.

[6]We note that θ can capture both natural recall and cues other than direct messages, e.g., general advertisements or news about catastrophes.

[7]Guilt could, however, also influence the weight put on warm-glow utility and hence giving behavior, but in the present paper we abstract from such effects to focus on the dynamic problem of whether to unsubscribe.

reminders induce the same cost.[8]

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 $k \in \{1, 2, ..., T-t\}$.[9] It follows 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

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

where $0 < \delta < 1$ is the inter-temporal discount factor and $E[\cdot \mid \Omega_t]$ 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.[10]

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 finite horizon allows us to investigate the effect of varying the horizon. 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 thus allow us to test it.

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

[8] Clearly a constant annoyance cost is a simplification. There are other potential functional forms and heterogeneous interactions with donor types that might affect the size of the annoyance cost in each period and for each donor. Nevertheless, in respect to the realism-tractability trade-off we have decided that the constant approach is the most straight forward. Further work is needed to determine a preferable functional form.

[9] While the irreversability is stringent assumption, this is in line with our experimental setting. For technical reasons, once unsubscribed donors can only resubscribe with a new e-mail address.

[10] We think it is reasonable that the donor can both forget about a donation e-mail, but at the same time be aware of his past warm-glow utility in the same way as one might forget to send a friend a birthday card, but does not forget how one generally feels about this friend, although the precise emotion might vary over time.

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. One can think of this assumption as capturing a small cost of pressing the unsubscribe button that tips the balance in favor of not unsubscribing. Similarly, we assume that the donor does not give anything if he is indifferent between doing so or not.

The model solution has a sequential structure: the donor 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 solution has classic threshold properties: A donor with a sufficiently low realization of a_t unsubscribes, while a donor with a high realization of a_t makes a donation. 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 and the amount donated is weakly increasing in a_t.[11]

In our model, the donor unsubscribes if he expects future annoyance costs to be larger than the "the option value" of subscribing i.e. the warm-glow utility foregone by not being reminded in the future. This is satisfied when a_t is sufficiently low. We note that current utility is unaffected by the current period unsubscription choice because the current annoyance cost cannot be avoided.

3 Experimental design and testable predictions

To test our model 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 and regardless of the length of a period.

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

3.1 Experiment I: A targeted reminder

Experiment I test the prediction of our model 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:[12] **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.

We note that although the targeted reminder is not explicitly announced, it is likely that donors in both treatments believe that they will indeed receive further messages in the future as assumed in our model. Hence applying the notation of the model to the experimental setting, we send an initial e-mail in period $t = 1$ in both treatments and an additional e-mail in period $t = 2$ to potential donors in the Targeted Reminder 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 Targeted Reminder than in Control I:* $P(g_1^{TR} + g_2^{TR} > 0) > P(g_1^{CI} + g_2^{CI} > 0)$. *And in particular in period* $t = 2$*:* $P(g_2^{TR} > 0) > P(g_2^{CI} > 0)$.

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

Intuitively, donors in the two treatments are equally likely to give and unsubscribe in the first period but donors in the Targeted Reminder treatment are more likely to give and unsubscribe in period $t = 2$ *because* they receive the second message prompting them to think about the donation and unsubscription decisions for the period $t = 2$ realization of a_t (the weight on warm-glow).

[12] Two further treatments, testing the effect of deadlines on donations, were implemented in parallel and are described in detail in Damgaard and Gravert (2017). An early working paper version (Damgaard and Gravert, 2014) also included data from the reminder treatment presented in this paper, but this data was excluded from Damgaard and Gravert (2017).

3.2 Experiment II: Changing the option value of subscribing

Experiment II tests the prediction of the model that unsubscription choices are affected by the option value of subscribing and beliefs about future annoyance costs. 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:[13] **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 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 predictions for Experiment II, let one period correspond to one month. Then $p_t = 1$ for all t in Control II and Future Benefit, 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., $d(g_2, a_2) = a_2 v(g_2) - c(g_2)$.

[13] We 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.

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

The mechanism is as follows: The match m reduces the cost of achieving a certain level of warm-glow utility for donations above the threshold X, i.e. $d_m(g_2^*, a_2) \geq d(g_2^*, a_2)$. Hence for a given value of a_1, the expected option value of remaining subscribed (at least until next period) is greater in the Future Benefit treatment than in Control II.

Donors in the Low Frequency treatment are told in period $t = 1$ 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 reduces the annoyance costs, but also the likelihood of being reminded to give. For donors with a low realization of a_t who are on the margin of unsubscribing, the effect of lower annoyance outweighs the effect of foregoing opportunities to donate if the probability of giving is sufficiently low.[14] Hence we have the following prediction.

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

In terms of giving behavior, the model predicts no differences across treatments, as giving is not constrained by an inter-temporal budget constraint. Effectively, Prediction 5 is therefore a test of the validity of this assumption.

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

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

[14] In a setting as ours with no social pressure costs, it seems reasonable to assume that the probability of giving, i.e. $P(a_t > \bar{a})$ is small. In Control I we find very little giving without asking.

in Denmark with a total revenue of 0.57 billion DKK in 2013 (DanChurch Aid, 2013). The total annual revenue of charities in Denmark is estimated to be about 2 billion DKK.[15] 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. Warm list means that these donors have donated previously to this 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 attention considerations. It also does not include high stakes donors. This leaves us with donors who infrequently donate amounts smaller than an hourly wage of the average Dane.

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

[15] Deloitte 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).

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[16], and the e-mail solicited money for poor children in Africa (see Figure J1 in the Online Appendix for a screenshot).[17]

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 J3 in the Online Appendix 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 entitle them to any ownership or rights regarding the store, i.e., it was a pure donation. 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 Online Appendix D.

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 and the change in solicitation cause, 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

Potential donors could unsubscribe from the mailing list by clicking on a link at the bottom of every e-mail. The design and visibility of the link was identical in

[16] All experimental designs face a trade-off between precision of the information and a natural setting. We focused on a completely natural setting, indirectly measuring annoyance costs by observing unsubscriptions. A less natural, but more direct way to measure annoyance costs is by eliciting the willingness to pay for a reminder as in Allcott and Kessler (2015).

[17] All experimental material and translations thereof is in the Online Appendix.

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	0.63	0.63	0.64	0.63
	(0.49)	(0.48)	(0.48)	(0.48)	(0.48)
Age (years)	46	46	38	38	38
	(15)	(15)	(15)	(15)	(15)
City (share)	0.33	0.33	0.35	0.36	0.35
	(0.47)	(0.47)	(0.48)	(0.48)	(0.48)
Amount donated last time (DKK)	300	313	191	194	192
	(622)	(553)	(408)	(419)	(516)
Months since last donation	35	35	32	32	31
	(19)	(19)	(22)	(22)	(22)
Months on e-mail list	1	1	24	24	24
	(-)	(-)	(5.5)	(5.5)	(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.

all e-mails and the unsubscription link was less salient than the donation button. If donors clicked "unsubscribe", 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. Choosing an answer and clicking on the unsubscribe button led to a confirmation message and a link to the homepage of the charity. A discussion of the survey findings can be found in Online Appendix B.

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 %)							
Perc. gave	0.44	0.35	0.53	0.66	0.65	0.67	0.65
	(6.6)	(5.9)	(7.3)	(8.1)	(8.0)	(8.2)	(8.0)
Perc. unsubscribed	2.90	2.14	3.67	0.38	0.49	0.30	0.36
	(16.8)	(14.5)	(18.8)	(6.9)	(7.0)	(5.5)	(5.9)
Observations (N)							
Full sample	17,391	8,692	8,699	43,489	14,501	14,494	14,494
# who gave	76	30	46	285	94	97	94
# unsubscribers	504	186	318	167	71	44	52

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

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.

5.1 Experiment I: The effect of a targeted reminder

Our model provides predictions regarding the giving and unsubscription behavior in Experiment I. We discuss the evidence for each of the two predictions in turn.

5.1.1 The reminder increases the number of donations

Figure 1 shows that the share of donors with positive donations in the Targeted Reminder treatment is larger than that in Control I (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) and we thus replicate the findings of Huck and Rasul (2010) that reminders can increase donations on the extensive margin. Table 3 provides the results of probit regressions on the likelihood of donating with and without controls.[18] The treatment effect is similar in sign and magnitude when including individual specific controls but is then insignificant. The number of months that have passed since the last donation negatively affect the probability of giving. This is consistent with the prediction of the model that donors have a high realization of a and suggests that people who have not donated recently are more marginal and therefore less likely to give now. On the intensive margin, we see a slight increase in the amount donated, conditional on donating, for people in the Targeted Reminder treatment compared to Control I as shown in Figure 1. 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 Online Appendix).

Figure 1: Giving and unsubscription behavior in Experiment I

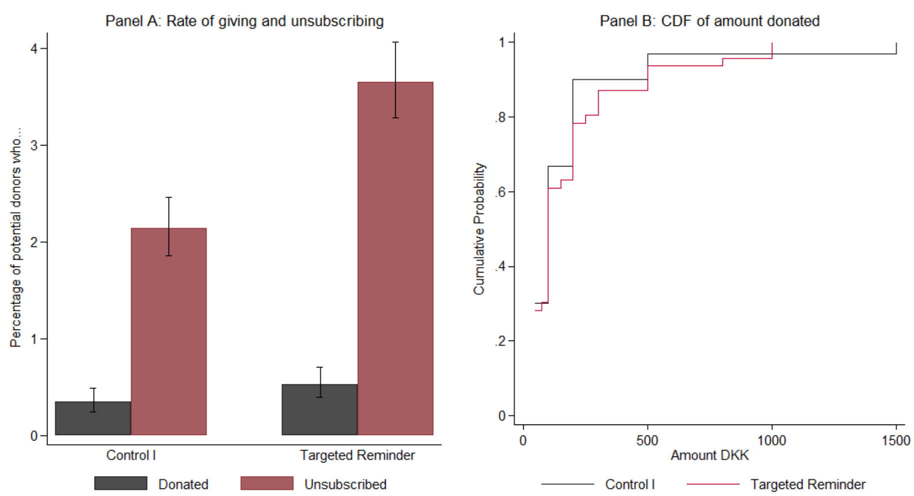

Notes: Panel A shows the rate of giving and unsubscribing with confidence intervals. Panel B shows the cumulative distribution function (CDF) of the amount donated conditional on giving. The difference in giving is significant at 10% level (Pearson chi2(1) = 3.3703, p-value = 0.066, Fisher's exact = 0.084) and the difference in unsubscription is significant at 1% level (Pearson chi2(1) = 35.4939, p-value = 0.000, Fisher's exact = 0.000). The differences in the CDF are not significant using a Mann-Whitney test, a two-sided two-sample t-test or a Kolmogorov-Smirnov test.

[18]We provide a robustness check for multiple hypothesis testing based on List et al. (2016) in Online Appendix E for all our estimation outcomes. The estimates are also robust to reestimation on the small sample for which all covariates are available.

Table 3: Donation and Unsubscription decisions in Experiment I and II

	Donated				**Unsubscribed**			
	Experiment I		**Experiment II**		**Experiment I**		**Experiment II**	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Targeted Reminder	0.00184*	0.00269			0.01516***	0.01114***		
	(0.00100)	(0.00195)			(0.00254)	(0.00398)		
Low Frequency			0.00021	0.00006			-0.00186**	-0.00192**
			(0.00095)	(0.00091)			(0.00074)	(0.00081)
Future Benefit			0.00000	-0.00006			-0.00131*	-0.00039
			(0.00095)	(0.0009)			(0.00076)	(0.00091)
Female		0.00131		0.00198***		-0.01804***		0.00027
		(0.00193)		(0.00074)		(0.00444)		(0.00083)
Age		0.00015**		0.00026***		-0.00035**		-0.00003
		(0.00007)		(0.00002)		(0.00014)		(0.00003)
City		0.00058		0.00319***		0.00645		-0.00001
		(0.00214)		(0.00093)		(0.00445)		(0.00086)
Months since last donated		-0.00027***		-0.00005**		-0.00022*		-0.00001
		(0.00006)		(0.00002)		(0.00011)		(0.00002)
Amount last donated		-0.00000		0.00000***		0.00000		-0.00000
		(0.00000)		(0.00005)		(0.00000)		(0.00000)
Months on e-mail list				-0.00018***				-0.00015**
				(0.00000)				(0.00007)
Observations	17,391	6,448	43,592	27,220	17391	6448	43489	27053
Pseudo R^2	0.004	0.046	0.000	0.079	0.008	0.023	0.003	0.014

Notes: The table provides the marginal effects and standard errors in brackets of probit regressions on the binary donation and unsubscription choices. 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.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

We find further support for Prediction 1 when looking at the timing of the donations. According to Prediction 1 the probability of donating around the time of the reminder should be higher in the Targeted Reminder group than in Control I. Figure 2 shows that most donations are made either on the day or the day after the initial solicitation mail. The differences in the numbers who donate on the first two days in the two groups is not significant. We also see a spike in donations on the day of the reminder but as predicted only in the Targeted Reminder group. In addition, 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 unless they are reminded.

Figure 2: Timing of giving in Experiment I

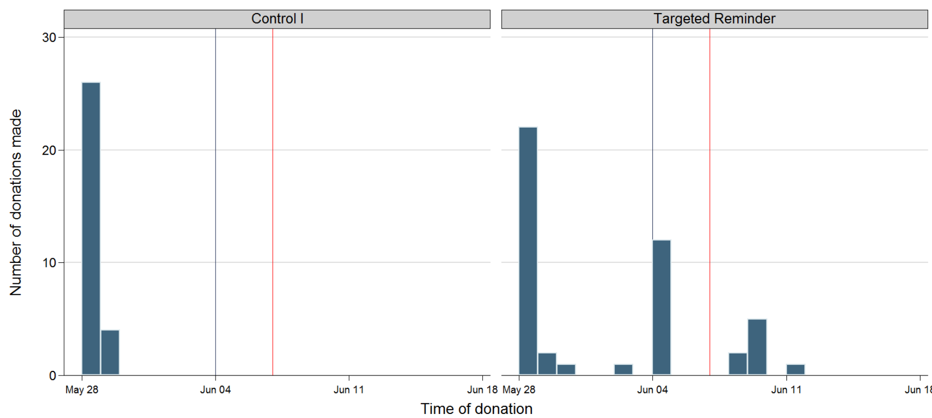

Notes: The figure shows the number of donations made on each day in Control I and Targeted Reminder. The initial e-mail was sent on May 28th, and the reminder on June 4th. The deadline for donating was June 7th. Data on the timing of unsubscriptions is not available for Experiment I.

5.1.2 The reminder increases the number of unsubscriptions

Figure 1 documents a large treatment difference in unsubscription behavior.[19] 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 Table 3, we see that the effect is robust to the inclusion of control variables. In-

[19] Data on unsubscriptions is only available for the treatment period for unsubscriptions through the links in the e-mails sent out as part of the experiment.

cluding all controls, the reminder increases the likelihood of unsubscribing by 1.1 percentage points compared to Control I.

5.2 Experiment II: Effects of changing the option value

The purpose of Experiment II is to test whether people (as implied by our model) account for the option value of subscribing when making their unsubscription decision. Compared to Control II, the treatments thus increase the value of remaining subscribed.

Figure 3: Giving and unsubscription behavior in Experiment II

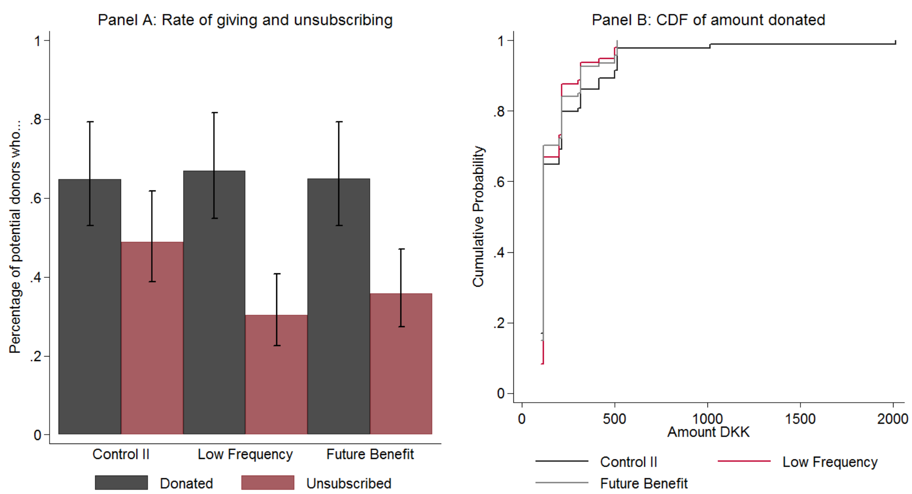

Notes: Panel A illustrates the rate of giving and unsubscribing with confidence intervals. Panel B shows the CDF of the amount donated conditional on giving. The differences in giving are not significant. The difference in unsubscription between Control II and Low Frequency is significant at 1% level (Pearson Chi2(1) = 6.35 (p-value = 0.01), between Control II and Future Benefit at the 10% level (Pearson chi2(1) = 2.94 (p-value=0.09)), and between Low Frequency and Future Benefit are not significant (Pearson chi2(1)=0.67 (p-value 0.41)). The differences in the CDF are not significant using a Mann-Whitney test, a two-sided two-sample t-test, or a Kolmogorov-Smirnov test.

5.2.1 A lower frequency of messages reduces the unsubscription rate

In line with Prediction 4, the Low Frequency treatment reduces the unsubscription rate from 0.49% to 0.30% (see Figure 3).[20] That is a reduction of 39%. The probit

[20] To measure the reaction in a clean way, we only analyze behavior within the first three days of sending the message. This is the time frame in which most unsubscriptions are carried out, and it

regressions in Table 3, columns 5-8 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 also goes in the direction implied by the model, i.e., Prediction 3. We find a marginally significant effect on the unsubscriptions compared to Control II, but this effect is not robust to the inclusion of controls.

The unsubscription rate is lower in Experiment II than in Experiment I. When we compare the unsubscription rates of our experiments with the rates the charity observed for some of their intermediate campaigns, we find that Experiment I is at the upper range of unsubscription rates and Experiment II at the lower range. Online Appendix Figure D1 shows a downward trend in the unsubscription rate over the past two years since Experiment I, with no visible difference between donation request e-mails and other newsletters sent by the charity. As the e-mails of Experiment I were some of the first e-mails the donors received, individuals who dislike newsletters in general may have reacted to those first e-mails and left the mailing list by the time we ran Experiment II. Nevertheless, there continues to be unsubscriptions every time the charity sends out an e-mail, and the unsubscribers are not just the most recent people joining the list (although being on the list for a longer significantly reduces the propensity to unsubscribe). The lower unsubscription rate in the second experiment does not seem to be explained by the slightly more difficult unsubscription process in Experiment II due to the attached survey question. We note that the number of donors who click on the unsubscription 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 to unsubscribe, this cannot explain the far lower unsubscription rate compared to Experiment I. Experiment I and II solicit donations for very different goods, which was already apparent in the subject line. Thus different sub-sets of donors might have opened the e-mails. Due to the temporal distance and different causes, the

helps reduce the noise created by other reminders or motivations for unsubscribing such as cleaning up the inbox after the 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.

experiments should be treated as independent evidence. Additionally, we find three people in each experiment who both unsubscribe and donate within a short time span (within 3 days). This is a low number compared to our total observations and could be due to reasons independent of the experiment such as closing an e-mail account.

For Experiment II, we have information on the timing of unsubscriptions (see Online Appendix Figure A1). Most unsubscriptions happen as an immediate reaction to the e-mail (within 3 days) and around 70 percent of all unsubscriptions happen on the day the e-mail is sent. There are no visible treatment differences.

5.2.2 The option value does not influence giving

The treatments have no significant effect on the decision to give as shown in Table 3. This is consistent with Prediction 5. We also find no significant effect of the treatments on the average amount donated, (see Online Appendix Table A1).

As in the first experiment, we find that most donations are made on the first two days after the e-mail is sent (see Online Appendix Figure A1). This shows that giving is either an immediate reaction to the solicitation or otherwise forgotten.

6 Structural estimation

The results of the two experiments are consistent with our model and provide evidence that reminders have a real welfare cost that people anticipate. Experiment I shows that nudging with reminders can have adverse effects and Experiment II shows that both the future value of warm glow and the expected annoyance cost affect unsubscribing as predicted by the model. In this and the following section we complement the reduced form analysis with a structural estimation of the model parameters (including the annoyance cost) which is then used as inputs for a welfare analysis. Clearly, assumptions have to be made. In addition, to the model assumptions made so far, we have to make some additional assumptions in order to be able to identify the model parameters. Therefore, the structural estimation should be taken as an approximation of the underlying parameters and not at face value.

The estimated version of our model is given by $v(g_{it}) = \log(1 + g_{it})$ and $c(g_{it}) = g_{it}$.[21] 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 let $p_t = 1$, meaning that the charity sends a message in every period, and we let a period correspond to one month.[22] To capture the lack of giving in some periods in the data where giving mostly appears to be reaction to a trigger from the charity, a natural disaster, or Christmas/the end of the tax year (see below), we set the probability of remembering without being reminded to zero, i.e., $\theta = 0$.[23] Finally, the monthly discount rate is calibrated to $\delta = 0.99835$ which corresponds to an annual real interest rate of 2%.

For this version of our model, $g_{it}^* = \operatorname{argmax}_{g_{it} \geq 0}(a_{it} \log(1 + g_{it}) - g_{it})$, which implies $g_{it}^* = a_{it} - 1$ for all $a_{it} > 1$ and $g_{it}^* = 0$ otherwise. Hence, 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 problem for which a closed form solution cannot be derived. Instead, we approximate the solution using backwards induction and evaluate conditional expectations by Monte Carlo integration.

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 had fixed the beliefs of the donors at the correct frequency. The data is restricted to individuals for whom historical donation data is available, giving a total of $N = 12,470$ individuals. We observe donation behavior for $\mathscr{T} = 54$ periods prior to the treat-

[21] The 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) and ensure that the estimate of Λ is measured in DKK (and in DellaVigna et al. (2012) that the social pressure cost is in Dollars).

[22] A 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.

[23] Alternative calibrations with $\theta > 0$ could capture overconfidence about prospective memory for donation decisions in the spirit of the findings by Ericson (2011). Such alternative calibrations scale the estimate of the annoyance cost by $1 - \theta$ because it reduces the anticipated value of being reminded by $1 - \theta$. All other estimates are unaffected.

ment 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 4 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.[24] 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. Figure 4 also shows the time series for the average amount donated conditional on donating.

Figure 4: Donation behavior over time

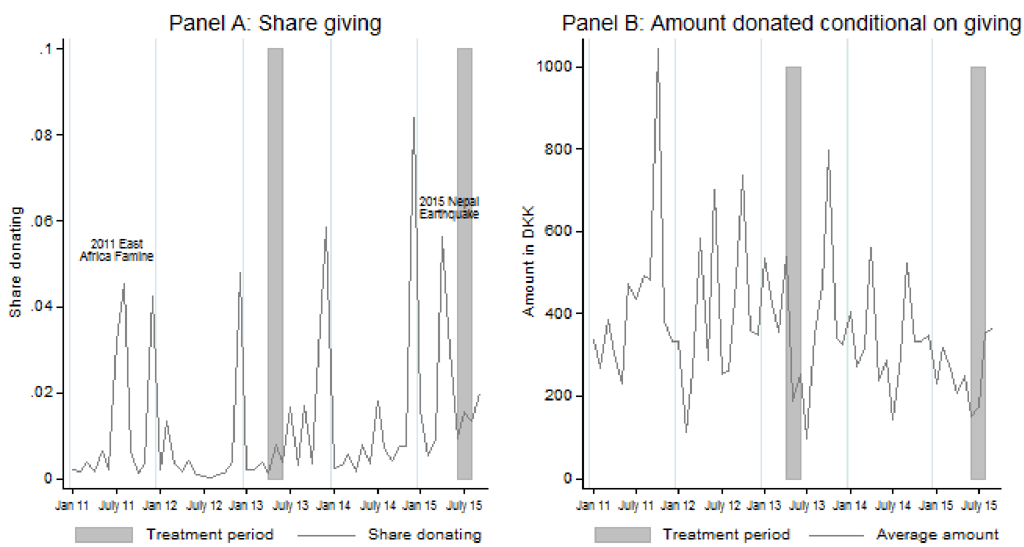

Notes: Data from Experiment II sample. Excludes payments made using Direct Debit and cash.

We decompose the estimation of our structural parameters. The parameters in the process for the weight on warm-glow can be identified solely from historical donation data independently of the donor's planning horizon T. The annoyance cost must be identified from unsubscription data and therefore depends on the planning horizon. We first estimate the warm-glow parameters using historical data prior to the treatment period, and then identify the annoyance costs from the unsubscription behavior in the treatment period given our estimate of the previous step. For the

[24] Cash 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.

estimation we use the method of simulated moments (MSM) following McFadden (1989).[25]

In the first step, we estimate μ, σ and ρ by considering the moments: i) the probability of not giving, ii) the probability of giving 100 DKK or less, iii) the probability of giving between 100 DKK and 200 DKK, iv) the probability of giving between 200 DKK and 400 DKK, v) the probability of giving between 400 DKK and 600 DKK, vi) a measure of the auto-covariance in giving $E(g_{it}g_{it-1})$.[26] The first five moments allow us to identify μ and σ in the distribution of a_{it} because of the simple relationship between the draw of a_{it} and giving in the version of the model considered here. Our last moment is required to pin down the persistence ρ in the process for a_{it}. In the estimation we minimize the weighted distance between the moments computed in the empirical data and the corresponding model moments computed by simulating donation data from the the model. For the model moments we simulate data for $\mathscr{T}$ periods across $\pi_1 N$ individuals, where $\pi_1 = 20$ is a scaling factor. The estimation is carried out using the standard procedure to obtain the optimal weighting matrix.

In the second step, we use the first-step estimates to impute starting values for a_{it} from the unconditional distribution, and determine Λ with the unsubscription rate serving as our only moment. Our setting is therefore just identified, making the weighting matrix redundant in this step. We compute the model-implied moment using Monte Carlo integration and determine Λ as the level required in the model in order for it to generate the same unsubscription rate as that observed in our data. When estimating the standard errors, we use standard procedures to account for the fact that Λ is estimated conditional on the estimates from the first step.

Our model matches the low propensity to give in the empirical data and the empirical distribution of giving in different ranges quite well (see Table 4). 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).

[25] The technical details are deferred to the Online Appendix which also explores the asymptotic properties of the MSM estimator in a Monte Carlo Study.

[26] To avoid perfect linearity between the chosen moments and hence insure full rank of the Jacobian, we do not include the probability of giving more than 600 DKK in the matched moments.

The unconditional mean of a_{it} is estimated to $\mu = -1,274$ and the standard deviation σ to 704. The negative mean captures the fact that on average across the pre-treatment period 98.79% of individuals do not donate in a given month. Hence, a large share of the distribution of a is below the donation threshold. At the same time the distribution of a must capture the tendency that people typically give between 100 and 200 DKK when they give. We note that the model implies a zero donation utility to potential donors 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 $\rho = 0.22$, suggesting that the warm-glow parameter is serially correlated.

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 ($\simeq$ \$1.95). 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. Calibrating different planing horizons, we find that the annoyance cost increases 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 ($T = 12$).[27] 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 potentially give. Hence, the potential donor will be relatively more likely to unsubscribe, and a relatively low annoyance cost is needed to explain the unsubscription rate in our data.

7 Welfare implications of reminders

One of the most difficult procedures in economics is to create an all encompassing welfare analysis of a policy. This section uses input from our structural estimation to give an indication of the potential welfare effects from reminders. Important

[27] See Online Appendix Figure I1 and I2 which provides estimates for different values of ρ.

Table 4: Structural estimates

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

Notes: The top 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., $N\pi$.

caveats are discussed.

Our approach is as follows: We use the structural estimates to approximate the utility of potential donors who receive of a regular e-mail reminding them to donate. Then we simulate donation paths allowing us to approximate counterfactual behavior for unsubscribing donors and analyze the welfare effects for the charity.

A potential donor who does not give in response to the reminder incurs the annoyance cost Λ without the warm-glow from giving. We use our estimate of Λ, -12.95 DKK, to approximate the negative welfare effect. A potential donor who gives in response to the reminder experiences warm-glow utility in addition to the annoyance cost. We estimate the average positive welfare effect for those who give to 1,191 DKK (see Table 5). On average 1.2% of donors give in a given month, implying that when accounting for the hidden costs of nudging, the average welfare effect for recipients is positive and equals 1.50 DKK.

Our results also suggest that by *not* accounting for the hidden costs of nudging, the welfare effect would be overestimated by a factor of almost ten. In addition, one would overlook that the positive welfare effect for people who are nudged to donate comes at the cost of a welfare loss to the vast majority of people contacted who are nudged but nevertheless prefer not to donate.

Table 5: Welfare effect of an average fundraising e-mail

	Accounting for the hidden costs			Not accounting for the hidden costs
Potential donors				
Welfare conditional on giving	1,191.02			1,203.97
Percentage of potential donors who give (%)	1.2			1.2
Welfare conditional on not giving	-12.95			0
Average welfare effect per contacted individual	*1.50*			*14.45*
Charity				
Charity discount rate (%)	2%	10%	20%	-
Immediate effect: Money raised per potential donor	3.07	3.07	3.07	3.07
Long-term effect: Money lost per potential donor	5.27	1.75	0.91	0
Net fundraiser effect per contacted individual	*-2.19*	*1.33*	*2.16*	*3.07*

Notes: Figures are in DKK unless otherwise stated. Welfare, percentage giving, immediate, and long-term effects are calculated given the estimated parameters. Long-term effects for the charity are assumed to last 576 months (given that donors in the sample are 38 years old and have a life expectancy of 86 years).

For the charity there is a positive immediate effect of sending a regular reminder which is equivalent to the amount raised. On average the charity raises 3.07 DKK per contacted individual. However, there is a long-term cost of lost future donations from unsubscribers because unsubscribing is an absorbing state, i.e., once people have unsubscribed they cannot rejoin the list with the same e-mail address.[28] Since the negative effect of unsubscribing for the charity occurs over future periods, the long-term effects depends on the charity's discount rate. It is not obvious how to calibrate the appropriate discount rate for the charity because it depends on factors such as the market rate and liquidity constraints. We therefore show the long-term

[28]The literature on optimal catalog mailings (among others Simester et al. (2006); Gönül and Shi (1998); Gönül and Hofstede (2006)) has also advocated taking a long run perspective when contacting people. 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.

and net effects for the charity for annual discount rates of 2%, 10%, and 20%. The net effect for the charity may be negative, -2.19 DKK, if the charity is patient and only discounts future losses with a low annual discount rate (2%). Higher rates lead to positive estimates. The extend to which *not* accounting for the hidden costs of nudging leads to overestimation of the net fundraiser effect therefore depends on the discount rate. The money raised by the charity must be used efficiently to generate positive welfare effects. This could be achieved if we assume 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 or as a donation to a less efficient charity (Singer, 2009). Importantly, we estimate the welfare effects on a sample of donors who donate irregularly and who usually donate small amounts. If the behavioral implications carry over to larger stakes donors, then the negative welfare effects for the charity would be much larger. Thus the effect size can be seen as a conservative estimate of the effect.

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, if unsubscribers increase donations to other charities after unsubscribing, this would give a positive welfare effect not accounted for in our calculations. However, additional negative effects might also arise if potential donors in general become less attentive to the content of messages and reminders when they receive more of them.

An additional caveat is that the estimates of the long-term effect of unsubscriptions are based on the assumption that people forget to donate unless they are reminded. However, in our field setting the charity communicates with potential donors through several communication channels and unsubscribing only means that donors switch off *one* communication channel, namely e-mails. Hence, it is possible that unsubscribers will be reminded to donate even after they have unsubscribed. An additional analysis in the Online Appendix shows that unsubscribers appear to be more marginal donors in terms of giving than subscribers, both before and after the treatment month.

Furthermore, there is considerable variation in the donation rate across periods. This has important policy implications, as the welfare of potential donors on average will be smaller than the estimates provided in Table 5 in periods with below average donation rates and higher in periods with above average donation rates. Hence, one remedy to reduce annoyance costs and unsubcriptions could be to target reminders to periods and fundraising campaigns expected to have above average response rates.

Similarly, rather than sending e-mails to every donor, the charity could provide donors with an opportunity to self-select into reminders. Inertia might, however, lead to very low sign-up rates and default enrollment might be preferred (Thaler and Sunstein, 2003). As is common in fundraising, our setting involves default enrollment as our sample of individuals had previously consented to the charity storing their e-mail address. A better solution than an opt-in could be to allow individuals to adjust the frequency of the newsletter. For example, some individuals might only want to be reminded to give around Christmas. 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 in time.[29] The increase in online giving implies better access to more data on giving patterns meaning that this could be a feasible strategy in the near future. On the other hand, natural catastrophes or emergencies do not follow a predicable pattern but could create an external shock to a person's a_{it}, which would make him willing to donate but this could be difficult to predict based on past donation behavior. Nevertheless, a more precisely targeted approach to fundraising would help balance the social welfare costs.

Lastly, we acknowledge that a model of heterogenous annoyance costs could also have been considered instead of our model of intertemporal and interpersonal heterogeneity in the weight on warm-glow utility a_{it}.We note that, a high level of warm-glow and an average level of annoyance leads to the same outcome as a reminder that is low in annoyance combined with average level of warm-glow. In our model, this variation is absorbed by a_{it} rather than by variation in annoyance costs. A model with varying annoyance costs could have been experimentally tested with more or less annoying reminders. This is an interesting avenue for further research.

[29] Allcott 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.

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 structurally estimate the annoyance costs and the utility of giving, thus tying in the reduced form results with our behavioral model. 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 to subscribers is about 12.95 DKK each month, and on average the welfare of a reminder to subscribers is very small (1.50 DKK). Failing to consider the hidden costs of nudging in a standard welfare analysis leads to an overestimation of the welfare effect by a factor of ten. Furthermore, when accounting for the long-term effects of unsubscriptions on giving, the net effect for the charity of sending a reminder is just 1.33 DKK with a discount rate of 10% and is negative for discount rates of 2%.

The model we develop and test experimentally in this paper could be extended to other settings were 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. Unfortunately, our data is not rich enough to estimate the optimal frequency of solicitations. Nevertheless, our model is the first inattention model that theoretically shows that there is a limit to the amount of reminders and thus provides a first step towards determining optimal frequencies.

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.

References

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 (2017). Avoiding the ask: A field experiment on altruism, empathy, and charitable giving. *125*(3), 625–653.

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.

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. NBER working paper No. 21406.

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 (2016). Effective reminders. forthcoming.

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*.

Damgaard, M. T. and C. Gravert (2014). Now or never! the effect of deadlines on charitable giving. *Aarhus University Economics Working Paper 2014-03*.

Damgaard, M. T. and C. Gravert (2017). Now or never! the effect of deadlines on charitable giving: Evidence from two natural field experiments. *Journal of Behavioral and Experimental Economics 66*, 78–87.

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.

Ericson, K. M. M. (2011). Forgetting we forget: Overconfidence and memory. *Journal of the European Economic Association 9*(1), 43–60.

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.

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 indsamlingsorganisationers indtægter mv. i perioden 2009 - 2013.

Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman (2016). Getting to the top of mind: How reminders increase saving. *Management Science 62*(12), 3393–3411.

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.

Letzler, R. and J. Tasoff (2014). Everyone believes in redemption: Nudges and overoptimism in costly task completion. *Journal of Economic Behavior & Organization 107*, 107–122.

List, J. A., A. M. Shaikh, and Y. Xu (2016, January). Multiple hypothesis testing in experimental economics. Working Paper 21875, National Bureau of Economic Research.

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*.

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 peternalism. *American Economic Review 93*(2), 175–179. Pappers 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 Americal Medical Informatics Association, 19*, 696–704.