Toward Understanding the Giving Process: Deciding to Give versus Giving*

James Andreoni
University of California, San Diego and NBER

Marta Serra-Garcia
University of California, San Diego and CESifo

Ann-Kathrin Koessler
Queensland University of Technology

December 11, 2015

PRELIMINARY AND INCOMPLETE

Abstract

We begin our study of altruism by making a natural observation that the joy of giving really has at least two parts. First is the utility of deciding to give, and second is the utility of seeing the decision realized. Such a distinction does not matter when deciding to give and giving occur in the same period, but when giving occurs over time this distinction can have important consequences by creating an a asynchronous relationship between the flow of utility and the flow of dollars. We present a simple theoretical model based on common models of discounting, and test the predictions experimentally. Consistent with our model, we find that donations increase when individuals commit to donate in the future. Without commitment, pledges to donate in the future are often reneged upon, leading to the absence of an increase in donations relative to donations that are immediate. We develop and test an intervention to reduce reneging: a thank-you note. Thanking individuals significantly increases the number of pledges that are converted to final donations.

JEL classification: D64, D90, C91.

Keywords: prosocial behavior, charitable giving, pledging, intertemporal choice.

*Andreoni: University of California, San Diego, Department of Economics, 9500 Gilman Drive, La Jolla, CA 92093 (andreoni@ucsd.edu); Serra-Garcia: University of California, San Diego, Rady School of Management, 9500 Gilman Drive, La Jolla, CA 92093; Koessler: Queensland University of Technology, 2 George Street, Brisbane, QLD 4000, Australia. We would like to thank the audiences at the Stanford Institute for Theoretical Economics, SPI Annual Conference, University of Pennsylvania and NHH Bergen for helpful comments. Andreoni would like to thank the National Science Foundation, grant SES-1427355, the Science of Philanthropy Initiative, and the John Templeton Foundation for financial support.

1 Introduction

In the past 30 years the tremendous growth of research on altruism and charitable giving has taught us a great deal about the chartable interaction.[1] In particular, we know that there are many steps and pieces that sum up to the charitable experience. First, there is the emotional pull of seeing a need, then there is the dance between the potential donor and the recipient, one is formulating the ask while the other is contemplating avoidance of that ask (e.g., Dana, Cain and Dawes, 2006; DellaVigna, List, Malmendier, 2011; Trachtman et al., 2015; Andreoni, Rao, and Trachtman, 2016), or potentially seeking the ask.

Donors are cultivated over time, and fundraising drives transpire over weeks and months. This is in strong contrast to the lion's share of research on charitable giving is conducted. Laboratory experiments are often conducted to look at an interaction that is resolved no more than an hour. Field experiments often look at single interventions into this process, or subtle variations in the wording of the ask, or by whom.

In this paper we take a humble step down a new path to look at the charitable interaction more wholly. At each step of the process, from stirring someone's heart to involving them as a long term donor, there would appear to be some opportunity for consumption on the part of the donor, whether that consumption comes in the form of imagining their potential impact, signaling to themselves and others their worth as humans, avoiding guilt, seeing a difference made in another's life, or feeling a greater part of an even greater community. Here, we ask whether looking at each stage separately, at different points in time, we can both understand the transitions between steps, and provide better understanding of how it works and how it can be incentivized to work even better.

We start with the simplest proof of this concept by looking at two points in time: The decision to give and the act of giving itself. These are nearly always studied as a single incident, yet they are often separated in time. This is done on grand scale by moguls such as Bill Gates, Mark Zuckerberg, and Warren Buffet, who pledge a lifetime of giving. But they are also done at the local synagogue where congregants are often called upon to

[1] For reviews, see Andreoni, 2006; Bekkers and Wiepking, 2010; List, 2011; and Andreoni and Payne, 2013.

publicly pledge their annual contributions. It happens when listeners to National Public Radio call in their $50 and $100 pledges. And it happens when parents enroll their children in parochial schools and agree to a major "gift" to the parish.

Our hypothesis in this paper is that individuals experience a flow of utility during at least three point of a charitable interaction: first, in making a promise or a pledge to give; second, in deciding to actually execute a gift; and third, in finally paying that gift. We study all three acts, but we separate them by time, so that each can be understood on its own. A person who is asked to give to charity today faces all three points compressed into a single brief interaction. A person who instead is asked today to make a binding decision about a gift to be made in one week faces a slightly different decision. The joy of deciding to give, and the personal or social rewards it brings, can at least partly be consumed at the time of the decision, while the remainder of the joy of giving along with all the pain of parting with one's money is delayed into the future. Next, consider a person who is asked to decide today to pledge to give money in the future, but the pledge is not binding. They must at some point before the deadline decide whether they will fulfill the pledge, and finally they must execute their decision to either the delight or disappointment of the charity.

The findings of our experiment, which are also strongly supported by a simple yet powerful theoretical model, reveal that people are significantly more likely to decide to give today when the payment is made in a week (42%) rather than immediately (33%). We also find that those who can pledge to give in a week are even more likely to say yes to that (67%). However, ultimate donations do not increase, despite the doubling of expressed intent to give. The difference between the rate of pledging and the rate of confirming pledges gave us an ideal opening for an intervention. How could we use the intervening week to increase the attachment to the pledge and reduce the degree of reneging?

We chose a simple and gentle manipulation that is often used in the fundraising world to cultivate repeat donors. For a random subset of those who made pledges to give, we sent a thank-you email within hours of the pledge. The email only expressed gratitude and anticipation of seeing them later when they would, hopefully, confirm their pledge. This simple manipulation, made hours after the pledge and a week before the confirmation was

due, increased donations by pledgers up to 44%. This provides further evidence suggesting that indeed there are different moments where donors gain utility along the process of giving, and that these partial flows of the giving utility are important and not immune to even subtle manipulations.

The next section of the paper presents background on related research, followed by a section with the theoretical model that underpins our research question. In the model a person experiences utility, positive or negative, from being asked, deciding to give, committing to give later, pledging to give later, and finally from completing their gift. We employ a unique modeling approach where the good feelings of giving obey something akin to a conservation of utility. In particular, we assume that all the elements of the joy of giving can be split apart, enjoyed at different times in the process and returned if ones changes her mind about giving. This neatly avoids the utility arbitrage that could come from changing one's mind, sometimes repeatedly. In our model, some utility may actually be sacrificed from reneging on a pledge and no one can be better off from changing her mind.

We follow this with a laboratory experiment conducted over two weeks. In the first week subjects are presented with an appeal to give to poor African families through GiveDirectly (www.GiveDirectly.org). They are randomly assigned on one of four conditions: they are asked to give $5 immediately (Give Now), commit to giving next week (Give Later), asked to pledge to give next week (Pledge), or given the choice between giving immediately or pledging (Pledge or Give Now). The results are remarkably consistent with our model.

We view this as the first attempt to disentangle the various components of the giving process and to successfully manipulate one of those aspects. We hope this is the first of many studies to look more systematically at the process of giving and how the parts, though separate, are not independent, and to try to gain a deeper and more integrated understanding of giving.

2 Background

A large majority of the theoretical and empirical literature on charitable giving has studied giving that occurs at the same moment as it is solicited. While this matches the behavior of charities that solicit immediate gifts, many charities ask individuals to pledge to give in the future.[2] A main contribution of our paper is to study the separation between the decision to give and the act of completing the gift.

Recently, several laboratory and field experiments on charitable giving have examined the effects of varying the deadline of a fundraising campaign and sending reminders to potential donors (e.g., Huck and Rasul, 2010; Damgaard and Gravert, 2014; Knowles and Servatka, 2015). Their findings suggest that reminders can increase donations, but that deadlines have no significant effects on giving. A new strand of the literature has studied the impact of time pressure and cognitive load on giving (e.g., Rand et al., 2012; Tinghoeg et al., 2013; Kessler and Meier, 2014; Recalde, Riedl and Vesterlund, 2014). While initial findings suggested that fast decisions are more generous, this finding has been subject to much discussion in followup studies. A central difference between these studies and our study is that in previous studies the decision to give and the payment of the gift occur simultaneously, and hence they do not distinguish between the decision to give and the payment of the gift.

The studies that are most closely related to ours are Breman (2011) and Dreber et al. (2014). Breman (2011) demonstrates in two large-scale field experiments that asking current donors to increase their donations is more effective when increases are scheduled in the future, rather than immediately. Dreber et al. (2014) present a dual-self model in which the short-term self is altruistic, while the long-term self is selfish. They present results from an online experiment in which individuals played a dictator game that was either implemented

[2] An example is offered by National Public Radio, whose fundraising relied for years on donors to phone in during pledge drives to promise a donation that could be made by credit card over the phone, or by bill to be paid by later, if the donor agreed to confirm the pledge. The option to pledge was also offered to MBA alumni in a fundraising campaign by the University of Chicago's Booth School of Business (Baran, Sapienza and Zingales, 2010). Finally, a prominent initiative around pledging is "The Giving Pledge", a campaign championed by Warren Buffet to encourage the wealthiest people in the world to pledge to donate most of their wealth to philanthropic causes.

immediately or with a delay. Their findings reveal that individuals are less altruistic when the dictator game is implemented in the future (see also Kovarik, 2009). Compared to these studies, our paper focuses on charitable donations among new donors, exploring both situations in which there is a commitment to give in the future or only a pledge. Further, our findings differ from those in Dreber et al. (2014), charitable donations increase when they are paid for in the future, which may be explained by a disutility from sharing in a dictator game.[3]

3 Theoretical Model of Intertemporal Altruism

In this section we present a model of intertemporal altruism. The main innovation of the model will be our assumptions about when utility flows. In particular, when an act of giving transpires over several stages of decisions, perhaps revised decisions, and finally ending in (potentially) a transfer of cash, a person can disassemble the full joy-of-giving into several intermediate stages. At each stage, pieces of the utility that would flow in a single period decision are instead broken apart and spread amongst the different points in time.

We model a situation in which someone has been asked to make a particular charitable gift of g, to which they can say yes or no. At first we will discuss the case where the individual has been asked to pay for the donation right away. We will then analyze the case where the individual is asked to pay for the donation in the future, with and without commitment. We will therefore introduce a separation between the time at which the individual decides to give and the moment at which the gift is transacted, which will allow utility from giving to be experienced at different stages. This will generate new insights regarding dynamics of giving.

Overall, we will examine the decision to donate to charity under four specific circumstances. These circumstances reflect the conditions we consider in our experimental analysis, hence we will talk about all four in the same language we use to discuss our experi-

[3] As discussed in further detail below, sharing in a dictator game may be perceived as unpleasant (e.g., Dana, Cain and Dawes (2006) show that subjects are willing to pay to avoid the task). If so, this could explain the difference in dynamics between our experiment and theirs.

mental design in the following sections.

All of the results are derived within a simple framework, in which the individual derives a direct utility from the act of giving. However, the same results can also be generated in a reproduction of important models of social-image signaling (Andreoni and Bernheim, 2009) or self-image signaling (Bénabou and Tirole, 2005) when giving.[4] We view the simple approach here to be preferable, not simply because it is trivial to manipulate (although that is a major advantage), but we also feel that it is a natural way to capture the fundamentals of the cognition in such problems, and therefore be more flexible in its interpretations. In the end, our extremely sparse model is surprisingly powerful in its predictions.

3.1 Condition 1: Give Now

We begin with the well-known case without a future. Here, as in our experiment, someone has been asked to give immediately a particular charitable gift of g, to which they can say yes or no. Several previous studies have established that people dislike saying no when they are asked.[5] Hence, we make the assumption that saying no yields utility of $-n < 0$. We assume that the act of giving, however, yields positive utility. Since our gift is fixed at g we will let $w > 0$ be the utility from the act of giving and $-g$ the utility cost of the forgone dollars. Thus, saying yes yields a net utility of $w - g$.

Combining these, one will say yes to the request to give if

$$w - g \geq -n. \tag{1}$$

3.2 Condition 2: Give Later

Imagine the same situation as above, except the person is asked now, at time $t = 0$, to commit to giving at some point in the future, time $t = 1$. Again, w is the utility from giving. However, we break with the prior literature here in a simple, natural, and perhaps

[4]The exact details of the models and the ensuing results can be obtained from the authors.

[5]See Dana, Cain and Dawes (2006), Dana, Weber and Xuang (2007), Andreoni and Rao (2011), DellaVigna, List, and Malmendier (2012), and Andreoni, Rao, and Trachtman (2016), among others.

surprisingly powerful way.

Simple introspection tells us that, just as the act of deciding not to give may make us feel uncomfortable, the act of deciding to do something positive also brings its joys. Thus, a giver might enjoy saying yes in the same way he would despise saying no. There is, however, an important point of caution in this assumption. If we allow people to get joy from deciding to give, we need to prevent the creation of wellsprings of utility that would arise if people would constantly change their minds. Thus, we assume that the total utility flowing from giving (w) cannot be expanded by the simple act of making a decision to give, or from changing one's mind. Rather, the utility at the time of deciding must, by necessity, reduce the joy at the time of giving. For our purposes we will assume the tradeoff is one-for-one.

Specifically, at the time of deciding we will assume that saying yes yields benefit $y > 0$. Then, when the gift is actually transacted, the person will gain the utility remainder of w, but also pay the cost of giving. This yields utility at the time of giving of $w - y - g$.[6]

Imagine deciding today, $t = 0$, on a binding promise to give g in period $t = 1$. Let δ be the discounting parameter. Then, if the utility from saying yes today, plus the discounted utility of giving later exceeds the utility of saying no today, an individual will say yes today:

$$y + \delta(w - g - y) \geq -n. \tag{2}$$

Rearranging this we find a more convenient expression,

$$w - g \geq -n - (n + y)\frac{1 - \delta}{\delta}. \tag{3}$$

Since $\delta < 1$, the expression $-(n + y)(1 - \delta)/\delta < 0$. This means that any individuals satisfying (1) will also satisfy (3). Likewise, many who where unwilling to give when the gift had to be paid immediately will now satisfy equation (3) and would be willing to

[6]Someone deciding and giving all in one period may have the same flow of utility, but they all simply occur in close succession. Thus, the utility of giving now is $y + w - y - g = w - g$, which is the same as shown in the last subsection. If deciding and giving are not synchronous, however, there may be opportunities for this to affect decisions.

commit in time 0 to giving later. Thus, total giving in Give-Later will exceed total giving in Give-Now. This leads to the first prediction.

Prediction 1: *Donation rates in Give-Later will be higher than in Give-Now, since those who would be willing to give immediately, equation (1), now will also be willing to commit now to giving later, equation (3), while some of those who were unwilling to give now will be now be willing to commit today to giving later.*

Interestingly, the result does not depend on temptation or self-control, but only on allowing *some* of the utility from giving to be consumed at the moment of decision, while delaying the the rest until the moment of the transfer. Notice as well that the prediction here is exactly the opposite of that reached by Dreber, et al. (2014), in a dual-self model of self-control. Their model predicts that the farther out into the future the giving occurs, even though one commits to the gift today, the lower the gift.[7] This contrast will allow us to compare the two approaches.

3.3 Condition 3: Pledging

What if the pledge to give as just described in Give-Later were not binding? That is, a person could pledge to give today, and then renege when the time comes to confirm that pledge.

The opportunity for reneging has several interesting potential consequences. First, how much utility flows to one at the time of saying yes to a pledge to give likely differs from that flowing when one says yes to a commitment to give. Intuition would tell us that the pledge is less valuable, so that the utility from pledging should be bounded above by the utility of commitment. Let y_p be the utility from a pledge. Then, we would expect $y_p \leq y$.

A second consequence of reneging is that, to the extent that others had expected a dona-

[7]The only way our model creates the effect predicted by Dreber et al. (2014), is if giving is an unpleasant act and people instead choose to consume some of this negative utility in the present. Such an assumption may not be out of place for a laboratory Dictator game, however, where we have seen significant numbers of subjects willing to pay to avoid the task altogether (e.g. Dana, Cane, and Dawes, 2006). If the experimenter demand effects are high and/or subjects feel pressure internally to conform to a norm or binding ethic, then this could create giving that fully reduces utility.

tion, reneging on a pledge is a form of deceit, or perhaps may not be fully distinguishable from such. Several studies have shown, moreover, that breaking promises or lying about intentions generates utility costs to individuals.[8] It seems sensible, therefore, to allow for a potential utility cost for individuals who make pledges and then renege. Let us denote this cost of reneging $r \geq 0$.

Pledging to give thus means that a person can say "yes, probably" immediately, experiencing y_p. As before, if she confirms the pledge, in the future she will receive utility $w - g - y_p$. Instead, if she reneges she will receive $-n - y_p - r$. Thus, an individual confirms the pledge if $w - g - y_p \geq -n - y_p - r$, or $w - g \geq -n - r$.

To allow for reneging in our model, we must introduce some randomness to the future period. A natural way to do this would be to assume that the values of y_p, n, w and r shift between the moment the pledge is made and the moment it has to be fulfilled. What matters, however, is how these shift in relation to g, the cost of giving. Thus, to simplify notation, we will assume that y, n, w and r are constant and that the utility cost of giving in $t = 1$ is a random variable, say $\tilde{g}$, where $\tilde{g}$ is distributed $f(\tilde{g})$, $-\infty \leq \tilde{g} \leq \infty$. We will assume that the expected value of $\tilde{g}$ is g, the known cost today.

Individuals will confirm their pledges if $w - g \geq -n - r$, that is, if $g \leq w + n + r$. Then we can write the utility of pledging as

$$\begin{aligned} U_P =& y_p + \delta \left\{ \int_{-\infty}^{w+n+r} (w - g - y_p) f(g) dg + \int_{w+n+r}^{\infty} (-n - y_p - r) f(g) dg \right\} \\ =& y_p + \delta \left\{ \int_{-\infty}^{w+n+r} (w - y_p) f(g) - \int_{-\infty}^{w+n+r} g f(g) dg + \int_{w+n+r}^{\infty} (-n - y_p - r) f(g) dg \right\} \\ =& y_p + \delta \left\{ p(w - y_p) - g_P + (1 - p)(-n - y_p - r) \right\} \end{aligned} \tag{4}$$

where g_P is the conditional mean of $\tilde{g}$ and $p = F(w + n + r)$ is the probability of giving.

A person will choose to pledge if it is better than saying no today:

$$y_p + \delta \left\{ p(w - y_p) - g_P + (1 - p)(-n - y_p - r) \right\} > -n.$$

[8] See Ellingsen and Johannesson (2004), Gneezy (2005), Charness and Dufwenberg (2006) and Serra-Garcia et al. (2013).

Since $pg > g_p$, this expression can be shown to hold if

$$w - g \geq -n - (n + y_p)\frac{1-\delta}{p\delta} + r\frac{1-p}{p} \tag{5}$$

In order to make (5) comparable to (3) we add and subtract $-(n+y)(1-\delta)/\delta$ to the right hand side of (5). Doing this and simplifying yields this expression:

$$w - g \geq -n - (n + y)\frac{1-\delta}{\delta} + \left[(py - (1-p)n - y_p)\frac{1-\delta}{p\delta} + r\frac{1-p}{p}\right]. \tag{6}$$

If the bracketed term on the right side of (6) is less than zero, then there will be more people pledging than there were giving in the Give-Later conditions.

To help evaluate (6), we illustrate the solutions one of the sensible ways to define y_p. When people are certain to give after pledging, so $p = 1$, we already assumed they experience y. In the same spirit, a pledge is a promise to give probabilistically, so it makes sense that we weight y by that same degree of uncertainty, that is, we assume that $y_p = py$.[9]

Substituting this into (6) we can write that condition for which a person will pledge as

$$w - g \geq -n - (n + y)\frac{1-\delta}{\delta} + \left[\frac{(1-\delta)(1-p)}{p\delta}(-n) + r\frac{1-p}{p}\right].$$

More individuals will pledge than say yes to giving in Give-Later if the amount in square brackets is negative, i.e. if

$$r \leq \frac{1-\delta}{\delta}n. \tag{7}$$

Given this, how will actual donations in the Pledge condition differ from the Give-Later condition? Consider first the simplest scenario where $r = 0$, and there are no utility costs of reneging. We know from (**??**) that more individuals will pledge than in the Give-Later condition. However, conditional on pledging, we know that only individuals with $g \leq w+n$

[9]A similar definition that would have also gotten similar conclusions is to also allow a portion of $-r$ to come forward to period 0, that is $y_p = py + (1-p)\phi r$, there $\phi = y/w$ is the same portion of w that the person can bring forward.

will confirm the pledge. Therefore, in this case, pledging will lead to the same donations as in Give-Now.

Now consider the case where $r > 0$ and condition (7) is satisfied. A first consequence is that more individuals will pledge than give in Give-Later. A second consequence is that, conditional on pledging, more people will actually give. At first this may seem counter-intuitive, since higher r should discourage giving. Higher r does deter some pledging to give, but conditional on pledging a higher r makes it more costly to renege on that pledge, hence those who by chance experience higher costs of giving will now pay that cost rather than absorbing the higher cost of reneging. This is shown early in this subsection when we stated that donors will confirm pledges as long as $g \leq w + n + r$. Thus, relative to Give-Later, when $r = 0$ we will expect many more pledges but about the same number of donors. As r increases, we expect the number of pledges to fall but the number of donors conditional on pledging to rise. The net effect is unclear and will depend on the shape of $F(g)$. If this $F(g)$ is symmetric around it's mean, then it is likely that total donations will change little by making pledges rather than commitments.

What if r is so high condition (7) is not satisfied. Then we will observe fewer pledges than in the Give-Later condition and, a consequence, fewer donations. Based on the existing evidence on lying costs and costs of breaking promises, however, we anticipate that that the costs of reneging on a promise will positive but not large and that condition (7) will be satisfied. This leads to the second prediction.

Prediction 2: *If condition (7) is satisfied, offering individuals the opportunity to pledge will lead to an increase in individuals pledging to give in $t = 0$, compared to individuals who commit to giving immediately or a week later in $t = 0$. However, individuals will also renege on their pledges in $t = 1$, and therefore the amount of donations is may not increase (if $r = 0$) or increase only slightly (with the costs of reneging, r).*

3.4 Condition 4: Pledge or Give Now

Imagine individuals are offered the opportunity to donate immediately, to say no immediately, or to pledge to give later. In this case, for those who pledge, the same considerations apply as discussed above. However, just as some enjoyed the flexibility offered in the Pledge condition, others will enjoy the commitment of the Give-Now condition.

Utility from pledging is as in (4). So a person in this condition will have utility

$$U = \max\{U_p, w - g, -n\} \tag{8}$$

corresponding to pledging, giving now, and saying no now.

If condition (7) is satisfied, we know that the U_p is greater than $-n$ for those who pledged in the Pledge condition, and so the identical set of people who will pledge in the Pledge condition will either pledge or give now in the Pledge-or-Give-Now condition. Stated differently, the fraction who say no today in the Pledge condition and in the Pledge-or-Give-Now condition will be identical.

Will any subject give now, and why? To take an extreme example, imagine a person for whom $n = 0$, that is, there is no cost in saying no and this person will surely give if she pledges. But this person will give now if $w - g > y$. It is easy to see, therefore, that those with higher values of w, p, $-n$ and r, that is, those who gain less by the delay, will all be more likely to move their giving up to the present.

What about final confirmed gifts? Since the set of potential donors is the same as in Pledge, and some of those will commit to giving in $t = 0$, the expected number of pledges will go up. However, since those who give now already had very high values of p in pledging, the increase in final donations may actually be small. This leads to Prediction 3.

Prediction 3: *If condition (7) is satisfied, the total number of people who in* $t = 0$ *pledge in the Pledge condition will be equal to the number of people who either pledge or give now in the Pledge-or-Give-Now condition. However, as some of those who can give now will do so, pledges in the Pledge-or-Give-Now condition will be lower. Expected final donations in Pledge-or-Give-Now will not be less than final donations in the Pledge condition, and*

could possibly be higher.

3.5 Stepping into the Gap with Thank-You Notes

Because our Pledge-or-Give-Now conditions separates by time the many components of the decision, and importantly also separates the utility experienced by these component parts, it is natural to think of using the gap between the utility flows to attempt to affect future choices. In particular, after the plegde to give has been made there is a gap until the final confirmation of that pledge is made. Is there a way we can influence remaining utility to enhance the likelihood of a pledge becoming a confirmed gift?

Looking to the world of charitable giving, a thank you note from the charity has often been found to be an effective tool for cultivating repeat givers. How would such a tool work in our model to convert more pledges into gifts? To have an effect, the thank you would have to come after the pledge, but before the confirmation of that pledge. How would a thank-you note change the confirmation decision?

Recall that the pledge will be confirmed if $w - g \geq -n - r$. This means there are four free parameters that could all be effected by the thank you, w, n, r and y_p. Intuitively, the thank you many make a person feel better about giving, thus increasing w. It may also make the person feel better about saying yes, thus increasing y_p, or more guilty if they end up saying no, thus increasing n or r. All of these effects of the thank-you note lead to the same conclusion.

Let Δ_w, Δ_n, and Δ_{y_p} be non-negative values that are added to w, n, and y after the thank-you. Since a thank you note has a very similar effect on the cost of saying no, as on reneging, let us set $\Delta_r = 0$. Now the pledge will be confirmed if $w + \Delta_w - \tilde{g} \geq -n - \Delta_n - y - \Delta_{y_p} - r$. Rearrange this, and we see the condition is $w - \tilde{g} \geq -n - y_p - r - \Delta$, where $\Delta = \Delta_w + \Delta_n + \Delta_y > 0$. Thus, for any realization of $\tilde{g}$, donations are more likely to be confirmed after the thank you note if any one of our free parameters is increased (or all are on on net increased) after the thank-you note.[10]

[10]Notice that it is also possible that the donor had given with the anticipation of a thank-you note. Thus, not getting a note would have the opposite effect: w, y and n would all fall and we could interpret Δ above

Prediction 4: *Individuals who pledge in the Pledge and Pledge-or-Give-Now conditions and subsequently receive a thank-you note will be more likely to convert their pledges to a gift than those who do not receive the thank-you note. Thus, the donation rate among those individuals who receive thank-you note in these conditions will be higher than the donation rate among those who did not.*

3.6 Empathy, Impulsivity, and Self-Control

A number of papers have recently uncovered interesting aspects of intertemporal altruism. In particular, if one sees an opportunity to give may be arriving in the future, people have been observed to avoid such opportunities (DellaVigna, List, Malmendier, 2012, Andreoni, Rao, & Trachtman, 2016, and Trachtman, et al., 2015). On the other extreme, when people are confronted with an opportunity to give and put under time pressure to decide, some researches have observed that people seem to more reflexively opt to give (Rand et al., 2012, 2014*a,* and 2014*b*), although many have also failed to find this effect (see Recalde, Riedl, & Vesterlund, 2014, for a review and discussion). Our model and analysis avoids both of these fault lines, but the model could inform them nonetheless.

First, we are describing a case that follows being asked. The model, however, allows that avoiding may be desirable by noting that someone who says no after being asked can actually be worse off than had the person never been asked.[11] In fact, we do not need to assume that giving yields positive utility to make giving desirable after being asked. Giving only requires that $w - g > -n$, but both sides of this equation could be negative. If so, then avoiding the ask would be the right choice and is indeed a moment for emotion-management as a means of self-control.

What about impulsive or reflexive giving? Our experiment differs in two ways from the literature on impulsivity and giving. First, we make a lengthy presentation about the charity,

as a negative number for those not receiving a thank you and 0 for those who do. Again, we would predict that those who are thanked will be more likely to confirm their pledges.

[11] More generally, assume utility from giving is produced with a function $w(g) - g$, which may be discontinuous at 0, but for all $g > 0$ we have $w(g)' > 0$ while $w(g) - g$ remains negative. Then we are considering a case where an individual has been asked to give g has to choose between a menu of all negative outcomes.

GiveDirectly, that gives subjects ample time to reflect, and we apply no time pressure. Second, our subjects give to a charity and do not participate in a game with other subjects. As such giving to charity can be seen as a public goods game, i.e. a positive sum game, that is often thought to yield extra psychic rewards from mutual cooperation. A dictator game, by contrast, is constant sum without any reciprocal altruism, and much less likely to have positive net utility than a public goods game. Indeed, public goods games more often support the hypothesis of impulsive giving than dictator games.[12]

Second, the financial costs of giving in our experiment are realized when they are received. The longer the delay in getting paid, the more likely one is to appear "impulsive." If subjects are from Mechanical Turk, they must go online to transfer credits from their payment account to an Amazon gift card, and if they are outside the US they are limited to one such transaction a day. Lab subjects, by contrast, typically have their payments in cash within a few minutes of their decisions. As expected, lab subjects are less likely to appear impulsive than Mechanical Turk subjects. While more investigation is needed, our model appears to capture the conflicting findings in this area.

3.7 Hypotheses

The predictions of our model lead to three central hypothesis, which we test experimentally. The first hypothesis concerns the decision at $t = 0$, while the second and third hypothesis regard the decisions made at $t = 1$.

Hypothesis 1 (Committing to giving at $t = 1$): *The fraction of people who give will be higher in the Give-Later condition than in the Give-Now condition.*

Hypothesis 2 (Pledging vs. Donations at $t = 0$): *The fraction of people who pledge to give in Pledge and Pledge-or-Give-Now conditions will be higher than the fraction of people who give in the Give-Now condition.*

Hypothesis 3 (Reneging): *Individuals who pledge in Pledge and Pledge-or-Give-Now and*

[12] On dictator game tests, see Piovesan and Wengström (2008) and Fiedler, Glöckner, Nicklish and Dickhert (2013).

receive a thank-you note will be less likely to renege than individuals who do not receive a thank-you note.

Hypothesis 4 (Final Donations): *Compared to Give-Now, final donations will increase in Give-Later, and in Pledge and Pledge-or-Give-Now when individuals receive a thank-you note. Without thank-you notes, final donations will increase only slightly in Pledge and Pledge-or-Give-Now compared to Give-Now.*

Next we present the details of the experimental design and then turn to our findings regarding these three hypotheses.

4 Experimental Design

The experiment consisted of two sessions, $t = 0$ and $t = 1$ sessions. The sessions were always spread exactly one week apart from each other.

In $t = 0$ of the experiment subjects were presented with the opportunity to donate \$5 to GiveDirectly, a charitable organization. The work of the charitable organization was presented at the beginning of the session and subjects were asked for a donation out of their experimental payment. The timing of the donation varied across treatments. In the Give-Now treatment, subjects were asked to donate \$5 today. In the Give-Later treatment, subjects were asked to commit to donating \$5 in $t = 1$. In the Pledge treatment, subjects were asked to donate \$5 in $t = 1$, subject to final confirmation in $t = 1$. In the Pledge-or-Give-Now treatment, subjects were asked whether they would like to donate \$5 today, or \$5 in $t = 1$, subject to final confirmation.

More specifically, the pledge was formulated as follows: "Yes, I'd like to donate \$5 next week. Ask me again next week and I will make my final decision." We chose to formulate a pledge as a statement of an intention to give, subject to final confirmation for several reasons. First, we chose not to use the word "pledge", as it can be interpreted differently across individuals, that is, while some may interpret it as non-binding promise, others may interpret it as a commitment to give. Second, to make clear that the decision was not final,

we explicitly wrote that subjects would be making their final decision in $t = 1$. Further, by formulating the statement in this way we kept the first sentence ("Yes, I'd like to donate $5 next week") the same as in the Give-Later treatment. Third, we chose to include a statement of an intention to give. Although it was not formulated as "I promise to contribute $5 next week", it may have been viewed by subjects as a promise (see Hanfling, 2008, for a philosophical argument). Thus, the effects of pledging measured in the experiment can be considered a lower bound of the effect of pledging more generally.

In $t = 1$ all subjects were shown their donation decision in $t = 0$. In the Pledge and Pledge-or-Give-Now treatments, if subjects had pledged to give in $t = 1$, they were asked to make their final decision.

4.1 Thank-You Note

Thank-you notes were sent randomly to subjects who pledged to give in $t = 1$, in the treatments Pledge and Pledge-or-Give-Now. The note was sent via e-mail by 5:00 p.m. on the same day of the session in $t = 0$. It is also important to note here that all subjects received an email 24 hours prior to their $t = 1$ session simply reminding them to attend.

To examine the sources of the effect behind the thank-you note we designed a "weak" and a "strong" version. In the "weak" thank-you note subjects were thanked for their participation and their decision to pledge. They were told that their contribution would make an important difference in the life of recipients. The note closed with a reminder of next week's session, where they could complete the study and collect their payments. The "strong" thank-you note had the same opening sentence. Instead of telling subjects about the general importance of their donation, the text emphasized that the donation would go to a family in Kenya "like this one", and a picture of a family was shown. Both notes closed with a reminder of next week's session. Due to limitations in the sample size, we compare the effect of "weak" versus "strong" thank-you note in the Pledge-or-Give-Now treatment only. In the Pledge treatment only thank-you notes in the weak version were sent.

To examine the impact of the thank-you note on attitudes toward the charity, we also sent the thank-you note to individuals who made a final decision to give in $t = 0$. More

specifically, the thank-you note in its weak version was sent to randomly among individuals who decided to give immediately in the Give-Now condition and Pledge-or-Give-Now, as well as individuals who decided to commit to giving later in Give-Later. Such a thank-you note cannot have an impact on donations, as these were already final, but it may impact attitudes towards the charity in the long run. We thus investigate its effect on attitudes collected in $t = 1$ at the end of Section 4.

4.2 Experimental Procedures

The experiment was conducted at UC San Diego. There were 501 participants: 86 in the Give-Now treatment, 82 in the Give-Later treatment, 118 in the Pledge treatment and 215 in the Pledge-or-Give-Now treatment. We purposely recruited more subjects in the latter two treatments to have enough observations when examining the effect of a "thank you" note on giving.

Subjects received a show-up fee of \$6 in $t = 0$ and a show-up fee of \$20 in $t = 1$. Participants were informed in the recruitment email about the two parts of the experiment. An important concern in an experiment that is spread over time is attrition after the first session. To counter this, we chose to concentrate the payment in $t = 1$, paying a show-up fee of \$20. This, in combination with the reminder email the day before the second session led to very high show up rates of 92% on average. Across treatments, there is an 95% show-up in the Give-Now treatment, 88% in the Pledge-or-Give-Now treatment, 93% in the Pledge treatment and 92% in the Pledge-or-Give-Now treatment (Chi-squared test, $p = 0.493$). The same experimenter – out of two who conducted the experiment – always ran week 1 and 2 sessions.[13]

We collected several additional measures of individual characteristics. In $t = 0$, we asked subjects to complete the Cognitive Reflection Test (CRT), developed by Frederick (2001). We also asked subjects for their gender, age, ethnicity, major, and fluency in English. In $t = 1$, we asked several questions regarding their impulsivity and empathy, as

[13]Experimenter fixed effects are included in our analysis below.

well as their perception of the experiment.[14]

5 Results

In what follows we first analyze decisions in $t = 0$. We then examine decisions in $t = 1$ and compare final donation rates across treatments. We end with an analysis of the relationship between empathy, impulsivity and giving.

5.1 Decisions at $t = 0$

Figure 1 displays donation decisions in $t = 0$. In the Give-Now treatment, 32% of the subjects choose to donate \$5. Giving increases to 43% in the Give-Later treatment. As shown in Table 1, the increase in giving is statistically significant ($Z = 2.31$, $p = 0.02$). The effect is also economically significant: we observe a 12 percentage point increase in the donation rate, which indicates a 37% increase relative to donations in the Give-Now treatment. This result reveals that allowing individuals to commit to donating later increases giving. It is consistent with our model of giving, whereby individuals derive a joy of giving at the moment the decision to give is made. In consequence, when payment of the gift is delayed into the future, the likelihood of giving increases.

Providing individuals with the opportunity of pledging leads to significant drop in the likelihood with which individuals will immediately refuse to donate to charity. In the Pledge treatment 65% of the subjects state an intention to donate. In the Pledge-or-Give-Now treatment, 21% of the subjects choose to donate immediately, while 47% state an intention to donate. The magnitude of the decrease in the frequency of immediate "No" replies is between 32 and 40 percentage points compared to the Give-Now treatment ($Z = 9.11$ in Pledge and 8.80 in Pledge-or-Give-Now, $p < 0.01$). The drop in immediate refusals is also significantly larger than the drop observed in the Give-Later treatment, as

[14]Precisely, students were asked questions from the Barrat Impulsiveness Scale (Barrat, 1959), a psychological measure for impulsivity and questions from the Interpersonal Reactivity Index (Davis, 1980), which measures different aspects of empathy. We also elicited students' general attitude towards charitable giving, the charity GiveDirectly for which donations were collected and their perception of the experiment.

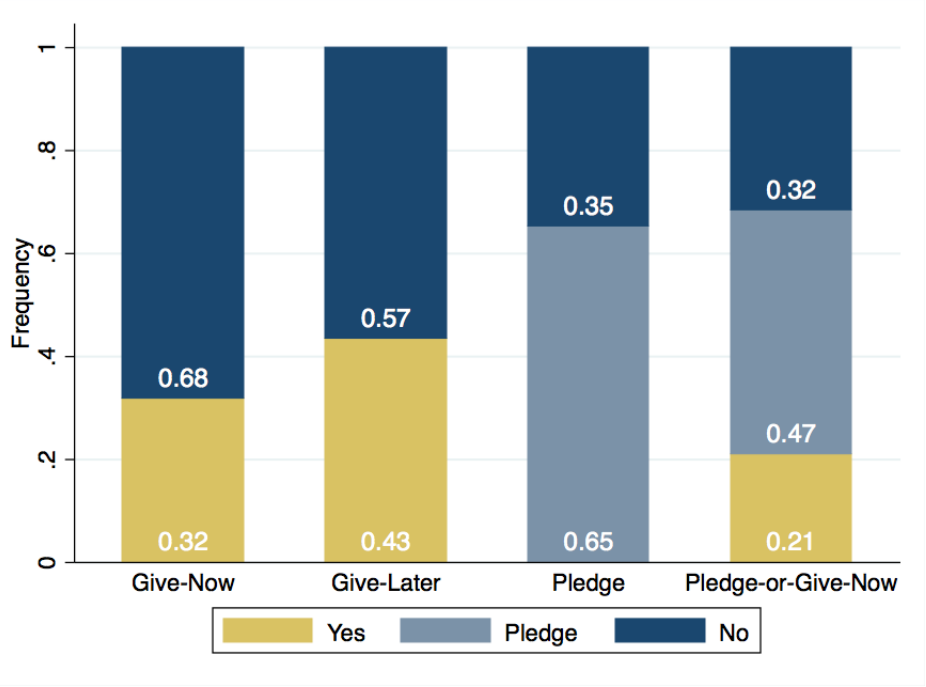

Figure 1: Donation decisions in $t = 0$

indicated by the Wald tests in Table 1. At the same time, no significant difference between Pledge and Pledge-or-Give-Now is observed.[15]

While pledging significantly decreases the frequency at which individuals immediately refuse to give to charity, there is a substantial proportion of individuals, between 32% and 35%, who choose to say "No" in $t = 0$. This consistent with the assumption in our model that, for some individuals, it is costly to pledge with the intention to renege on the pledge. Thus, individuals who anticipate they will renege on a pledge with a high probability prefer to say "No" immediately.

Overall, two results are obtained from decisions in $t = 0$:

Result 1: *Allowing individuals to commit to giving in the future, significantly increases giving.*

Result 2: *Allowing individuals to pledge, significantly decreases the likelihood with which*

[15]The results remain qualitatively similar if individual characteristics, such as age, gender, ethnicity, English-language skills and a measure of cognitive reflection (through the Cognitive Reflection Test by Frederick (2002)) are added as controls, as shown in column (2) of Table 1. Individual characteristics do not affect giving decisions in $t = 0$, except for a marginally positive effect of English-language skills on giving.

Table 1: Treatment effects in $t = 0$

	(1)	(2)
	Likelihood of not donating in $t = 0$ (choosing "No")	
Give-Later	-0.120**	-0.129**
	(0.052)	(0.051)
Pledge	-0.340***	-0.320***
	(0.037)	(0.039)
Pledge-or-Give-Now	-0.373***	-0.401***
	(0.042)	(0.045)
Wald-test:	*p-value*	*p-value*
Give-Later vs. Pledge	0.000	0.000
Give-Later vs. Pledge-or-Give-Now	0.000	0.000
Pledge vs. Pledge-or-Give-Now	0.388	0.109
Observations	501	500
Individual characteristics	No	Yes

Notes: This table presents the marginal effects (calculated at the means of all variables) from probit regressions on the likelihood of choosing not to donate in t=0 (selecting "No" in t=0). The variables Give Later, Pledge and Pledge-or-Give-Now are dummies that take value one in the corresponding treatment, zero otherwise. The regression model in column 2 also includes individual characteristics (asian ethnicity, gender, English-language skills, cognitive reflection score). Robust standard errors, clustered at the session level, are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

individuals immediately refuse to donate to charity.

Results 1 and 2 are in line with Hypotheses 1 and 2. The large decrease in immediate refusals to giving when pledging is possible suggests that allowing individuals to pledge is a promising venue to increase donations. However, an important question is how often individuals follow through with their pledges, a behavior we discuss next.

5.2 Decisions at $t = 1$

Figure 2 displays the frequency with which individuals renege on their pledge in $t = 1$. For this purpose, we focus on the treatments Pledge and Pledge-or-Give-Now.[16] We first

[16]In the Pledge treatment 77 subjects pledged to donate in $t = 0$. Among pledgers, 47 received a thank-you note in the weak version via email, while 30 received no thank-you note. In the Pledge-or-Give-Now treatment 102 subjects pledged to donate in $t = 0$. Among these subjects, 49 who received a thank-you note in the strong version via email, 27 who received the weak version of the thank-you note and 26 who received no note.

describe reneging in the absence of thank-you notes and then turn to the effect of thank-you notes on reneging.

A majority of the individuals who pledge in $t = 0$ renege on their pledge in $t = 1$. In the Pledge treatment, if subjects receive no thank-you note, 54% renege on their pledge. In the Pledge-or-Give-Now treatment, this proportion increases to 73%. The change is marginally significant, as shown in Table 2 column (4) ($Z = -1.79$, $p = 0.07$). The higher frequency of reneging in the Pledge-or-Give-Now treatment is consistent with selection. Individuals who preferred to give right away, selected into doing so in Pledge-or-Give-Now, while these individuals only had the option to pledge in the Pledge treatment. Therefore, those who decided to pledge in the Pledge-or-Give-Now treatment exhibited had higher likelihood of reneging in the first place.

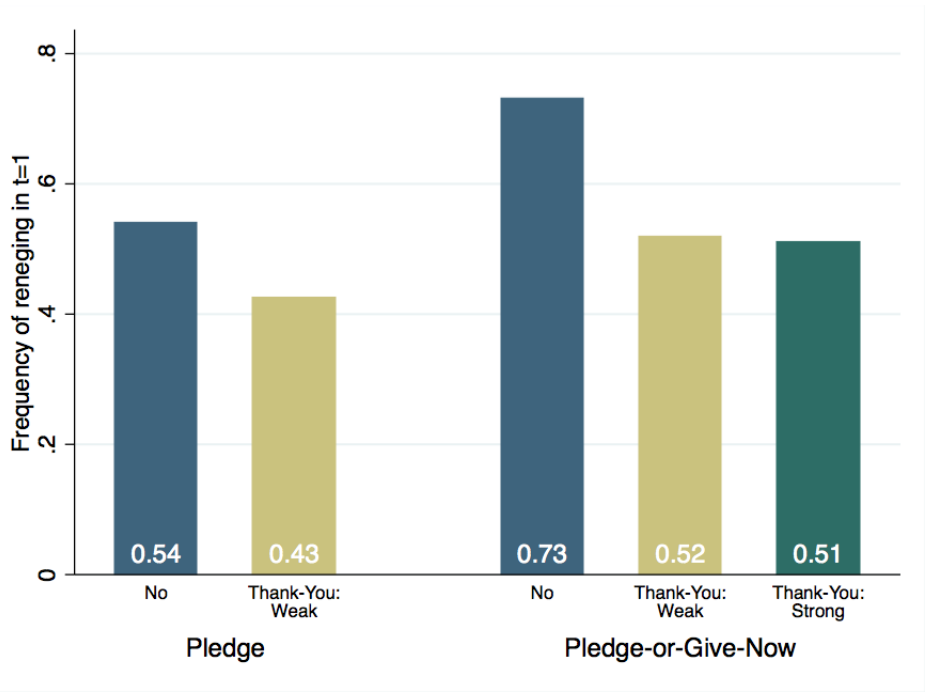

Figure 2: Reneging in $t = 1$

The thank-you note led to a decrease in reneging among pledgers. In the Pledge treatment, individuals who receive a weak thank-you note renege in 43% of the cases. The drop in reneging (from 54%) is however not significant, as shown in Table 2 column (1). In contrast, the effect of the thank-you note in the Pledge-or-Give-Now treatment is larger

and significant. Pledgers who receive a weak thank-you note renege 52% of the time, while those who receive a strong thank-you note renege 51% of the time. The drop in reneging is significant when considering both versions of the thank-you note jointly, as shown in Table 2 column (2), and marginally significant for each version of the thank-you note separately, as shown in Table 2 column (3).[17]

Table 2: Effect of thank-you note on reneging in $t = 1$

	(1)	(2)	(3)	(4)
	Likelihood of reneging in $t = 1$			
Treatment:	Pledge	Pledge-or-Give-Now		Pledge and Pledge-or-Give-Now
Thank-You	-0.081	-0.282**		-0.111
	(0.154)	(0.135)		(0.125)
Thank-You: Weak			-0.251*	
			(0.152)	
Thank-You: Strong			-0.314*	
			(0.165)	
Pledge-or-Give-Now				0.264*
				(0.148)
Pledge-or-Give-Now * Thank-You				-0.116
				(0.182)
Observations	76	102	102	178
Individual characteristics	Yes	Yes	Yes	Yes

Notes: This table presents the marginal effects (calculated at the means of all variables) from probit regressions on the likelihood of reneging on a pledge in $t = 1$. The variable Thank You is a dummy variable if the individual received a thank you note (weak or strong version), zero otherwise. The variables Thank You: Weak and Thank You: Strong are dummy variables that take value one if the thank you note was weak or strong, correspondingly. Pledge-or-Give-Now is a dummy that takes value one in the Pledge-or-Give-Now treatment, zero otherwise. All specifications include individual characteristics (asian ethnicity, gender, English-language skills, cognitive reflection score). Experimenter fixed effects included in all specifications. Robust standard errors, clustered at the session level, are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Result 2: *A majority of individuals renege on their pledge in $t = 1$. However, sending individuals a "thank you" note decreases the probability that the individual reneges on her pledge. The effect is significant in the Pledge-or-Give-Now treatment, implying that thank-you notes have the largest effect on individuals who have explicitly chosen not to donate*

[17]The same result from column (4) in Table 2. The effect of the thank-you note is not significant in the Pledge treatment (based on the coefficient of Thank-You). In contrast, it becomes significant in the Pledge-or-Give-Now treatment, when jointly testing the coefficients for Thank-you and Pledge-or-Give-Now*Thank-You ($\chi^2 = 3.00$, $p = 0.08$)

immediately, but to finalize their donation at a later point in time.

5.3 Donations

Donations made in the treatments Give-Now and Give-Later are final in $t = 0$. However, as we have seen, donations in treatments Pledge and Pledge-or-Give-Now are only final when they are confirmed in $t = 1$. Pooling all thank you conditions together, we observe that 34% of individuals donate in the Pledge treatment, while 41% donate in the Pledge-or-Give-Now treatment, as shown in Figure 3 Panel A. Donations increase significantly in the Pledge-or-Give-Now treatment, but they do not in the Pledge, relative to Give-Now, as shown in Table 3 columns (1) and (2).

An important question is what is the combined effect of offering the option to pledge and sending a thank-you note. In other words, can donations be increased when pledging is introduced together with thank-you notes? This requires examining the effect of the Pledge treatment combined with that of the thank-you note from an ex-ante perspective, i.e. considering all subjects and not only focusing on those who pledged to give, as we have done in Section 2.2. To measure this effect, it is important to note two facts. First, thank-you notes were assigned randomly across those individuals who pledged to donate in $t = 0$. Second, those who said "No" in $t = 0$ made a final decision, which could not be influenced by the thank-you note. Let us denote a final donation as $d \in \{0, 1\}$. The probability of observing a final donation conditional on having sent a thank-you note (in its weak version) in the Pledge condition can be written as follows: $P(d = 1|\text{Thank} - \text{you}) = P(\text{Pledge})P(d = 1|\text{Thank} - \text{you \& Pledge}) + P(\text{No})P(d = 1|\text{Thank} - \text{you \& No}) = 0.6525 \times 0.575 + 0.3475 \times 0 = 0.38$. Using the same conditional probability calculation for both the Pledge and Pledge-or-Give-Now treatments we obtain the results displayed in Figure 3 Panel B.

Figure 3 Panel B reveals that final donations in the Pledge and Pledge-or-Give-Now treatment differ by thank you condition. In the Pledge treatment, the likelihood of observing a final donation if an individual has not received a thank-you note is 0.3. When they receive a weak version of the thank-you note, this likelihood is 0.38. In the Pledge-or-

(a) Simple treatment effects

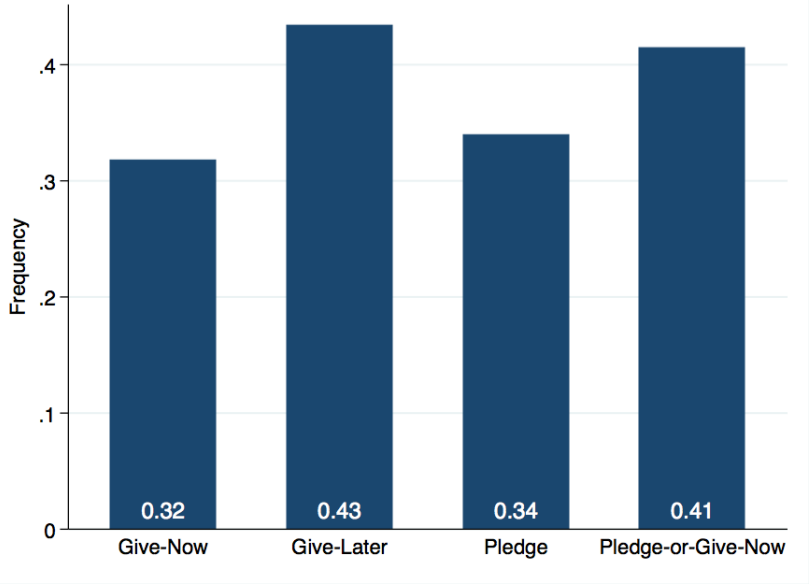

(b) Treatment and thank-you note effects

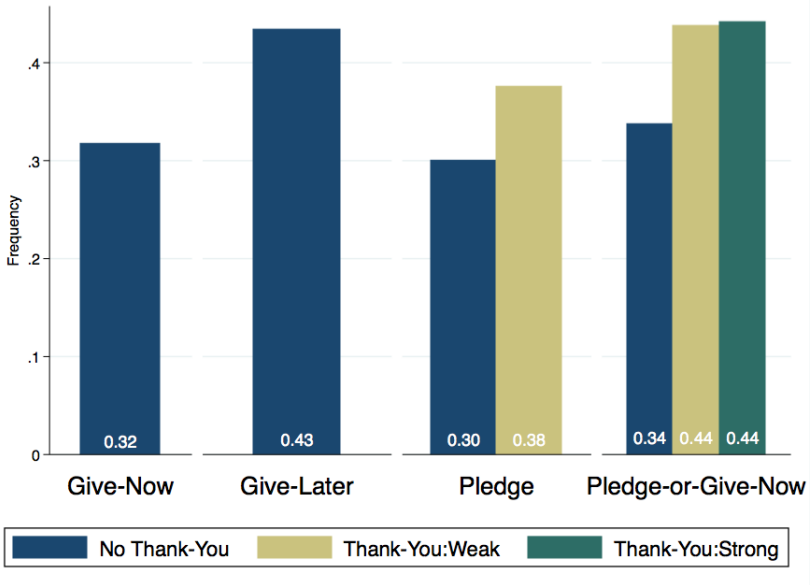

Figure 3: Final Donations

Give-Now treatment, the likelihood of observing a final donation if an individual has not received a thank-you note is 0.34, while it is 0.44 if she receives a thank-you note, in its weak or strong version.

To obtain significance tests based on donations by treatment and thank-you condition,

we use a bootstrapping method. In the experiment, individuals in a session were randomly assigned to receive a weak or strong version of the thank-you note (or no thank-you), conditional on pledging. Individuals who made a final decision in $t = 0$ were not assigned to a thank-you condition, as their decision was final. To examine the effect of pledging combined with thank-you notes, we randomly assign those individuals who made a final decision in $t = 0$ to one of the thank-you conditions. Intuitively, if 50% of individuals who pledge to donate in $t = 0$ received a weak thank-you note, 50% of the individuals who decided not to donate in $t = 0$ are randomly assigned to the weak thank-you note condition. To control for session and individual effects, we run a bootstrap analysis with 1,000 replications. We report the average regression coefficients and bootstrapped standard errors in Table 3 columns (3) and (4). We do not include the Give-Later treatment as there was no variation in thank-you conditions, but if we include it, results remain unchanged.[18]

The results in Table 3, columns (3) and (4), indicate that, without a thank you note, donations do not increase relative to the Give-Now condition. This result is in line with our model, by which no to small effects of pledging were expected without a thank-you note. By contrast, pledging combined with thank-you notes significantly increases donations. The average marginal effect of the weak version of the thank-you note is to increase donations by between 8.9 and 11.4 percentage points. The average marginal effect of the strong version of the thank you note is to increase donations by ca. 16 percentage points, relative to the Give-Now treatment. The difference between the effect of the weak and strong version of the thank-you note is not significant, though (χ^2-test, $p = 0.4244$ and $p = 0.3136$ for columns (2) and (3) respectively). This leads to Result 3.

Result 3: *Donations increase in the Give-Later treatment, and in the Pledge and Pledge-or-Give-Now treatments, if pledges are accompanied by a thank-you note.*

At the end of $t = 1$ sessions, we elicited subjects' attitudes towards the charity as

[18] In the Pledge treatment we aimed to send a thank-you note to 50% of the individuals, while in the Pledge-or-Give-Now condition we aimed to send a weak thank-you note to 25% of the individuals and a strong thank-you note to 50% of the individuals, in each session. The realized random assignment varied according to this distribution within each session. This randomness is also reflected in our analysis through the use of the bootstrap method.

Table 3: Treatment effects on final donations

	(1)	(2)	(3)	(4)
			Donation	
Give-Later	0.117**	0.128***	-	-
	(0.050)	(0.049)		
Pledge	0.022	0.037		
	(0.054)	(0.051)		
Pledge + No Thank-You			-0.042**	-0.026
			(0.017)	(0.018)
Pledge + Thank-You:Weak			0.089***	0.098***
			(0.020)	(0.020)
Pledge-or-Give-Now	0.097***	0.097**		
	(0.035)	(0.043)		
Pledge-or-Give-Now + No Thank-You			-0.043	-0.046
			(0.043)	(0.043)
Pledge-or-Give-Now + Thank-You:Weak			0.114***	0.109***
			(0.040)	(0.041)
Pledge-or-Give-Now + Thank-You:Strong			0.160***	0.169***
			(0.024)	(0.029)
Observations	501	500	418	417
Individual characteristics	No	Yes	No	Yes
Standard errors	Robust	Robust	Bootstrapped	Bootstrapped

Notes: This table presents the average marginal effects (calculated at the means of all variables) from probit regressions on final donation decisions. Columns 1-2 present the marginal effect from simple probit regressions on the treatment. Columns 3-4 present results from bootstrapped regressions, with 1000 repetitions. The variables Give-Later, Pledge, Pledge+No Thank-You, Pledge+Thank-You:Weak, Pledge-or-Give-Now, Pledge-or-Give-Now+No Thank-You, Pledge-or-Give-Now+Thank-You:Weak, Pledge-or-Give-Now+Thank-You:Strong are dummies that take value one in the corresponding treatment or treatment+thank you condition, zero otherwise. Columns 2 and 4 also include individual characteristics (asian ethnicity, gender, English-language skills, cognitive reflection score), and experimenter fixed effects. Robust standard errors, clustered at the session level, were used in each individual regression. Columns 1-2 display regular robust standard errors in parentheses, columns 3-4 present bootstrapped standard errors in parentheses. Normal-based confidence intervals are used to determine significance, but results remain qualitatively the same with percentile and bias-corrected confidence intervals. ***,**,* indicates significance at the 1%, 5% and 10% levels, respectively.

well as their perceptions about their decisions in the experiment (e.g., whether they regretted their donation decision). The aim of these non-incentivized questions was to examine whether there are spillover effects of the treatments and the thank-you notes on subjects' attitudes towards the charity and their own decisions. As we detail in Appendix B, the thank you notes had no systematic effects on subjects' attitudes towards the charity or their own perceptions about their donation decisions. We do not observe systematic effects of the treatments on attitudes either. A potential explanation is that these questions were not incentivized and hence cannot precisely measure spillover effects. However, the data sug-

gests that such effects, if present, were not large. Further, at the end of $t = 1$ sessions we elicited subjects' willingness to sign up to the charity's newsletter. There were no systematic spillover effects of delay and commitment on this decision either. Overall, these results suggest that introducing delay and changing commitment changes individual's incentives to donate, but does not directly affect the image of the charity. Such a result can be especially relevant for practitioners, who do not only worry about increasing donations in the short run, but also about potential long-term effects of any intervention.

5.4 Empathy and Impulsivity

Motivated by the growing literature on giving and impulsivity, we elicited survey measures of impulsivity and examine their relationship to giving decisions. In our context, where decisions are made without time pressure and cognitive load, we would expect such correlations to be weak, but strongly suggestive, if existent. At the same time, we elicited a measure of empathy with the aim of examining whether empathy, which could be part of the utility of giving (w), is related to giving decisions.

To elicit empathy we used the Interpersonal Reactivity Index (IRI) developed by Davis (1980), which measures empathy as the "reactions of one individual to the observed experiences of another" (Davis, 1983). To elicit impulsivity we used Barratt's Impulsiveness Scale (BIS; Barratt, 1959) and the Cognitive Reflection Test (CRT; Frederik, 2005).[19] In the BIS impulsiveness is defined as "as a predisposition toward rapid, unplanned reactions to internal or external stimuli without regard to the negative consequences of these reactions to the impulsive individuals or to others" (Moeller et al., 2001). The BIS has been widely used in psychology to assess the personality construct of impulsiveness. The BIS differs from the CRT in that the latter is a behavioral measure of impulsivity while the former is a self-reported personality trait. Behavioral and personality measures are typically uncorrelated (Stanford et al., 2009), as is the case in our sample. The number of

[19]The IRI and BIS were elicited at the end of $t = 1$ sessions, since these were lengthy questionnaires that discuss feelings regarding others and impulsivity in everyday decisions. The CRT was elicited at the end of $t = 0$ sessions, since it only consists of three simple and short math questions, which do not refer to others or to impulsive behaviors, and hence we did not expect an influence in $t = 1$ behavior.

correct answers in the CRT and the BIS score are not significantly correlated (Spearman's $\rho = 0.0203$, $p = 0.6648$). At the same time, the average BIS total score in our sample is 59.51 (SD $= 9.01$), which is not significantly different from the BIS total score obtained in other studies (62.3 with SD $= 10.3$; Stanford et al., 2009).

We analyze the relationship between empathy and impulsivity and donation decisions. The results are presented in Table 4. Columns (1) focuses on the decision to not donate in $t = 0$, while column (2) examines final donation decisions. BIS and IRI scores are standardized.

Table 4: Empathy, Impulsivity and Donation Behavior

	(1) Likelihood of not donating in $t = 0$	(2) Donation (in $t = 1$)
Empathy:		
Interpersonal Reactivity Index	-0.095**	0.144***
	(0.037)	(0.035)
Impulsivity:		
Trait: Barratt Impulsiveness Scale	0.001	0.019
	(0.021)	(0.031)
Behavioral: Cognitive Reflection Score	-0.034	0.018
	(0.025)	(0.020)
Give-Later	-0.121*	0.124*
	(0.066)	(0.068)
Pledge	-0.320***	0.054
	(0.035)	(0.050)
Pledge-or-Give-Now	-0.400***	0.105**
	(0.042)	(0.042)
Observations	459	459

Notes: This table presents marginal effects from probit regressions on the decision not to donate in $t = 0$ (column (1)) and donation decisions (column (2)) over the two time horizons. Empathy (IRI) is the standardized mean score in the Interpersonal Reactivity Index (Davis, 1980). Impulsivity is measured using the trait impulsivity measure developed by Barratt (1959) though the Barratt Impulsiveness Scale (BIS) and a behavioral impulsivity measure, the Cognitive Reflection Test (Frederick, 2002). The Baratt Impulsiveness Score is measured as the standardized total score of the Barratt Impulsiveness Scale (Stanford et al., 2009). The Cognitive Reflection Score is the total number of correct answers obtained in the cognitive reflection test. The variables Give-Later, Pledge and Pledge-or-Give-Now are dummies that take value one in the corresponding treatment, zero otherwise. All specifications include individual characteristics (asian, gender and native english speaker) as defined in Table 1. Experimenter fixed effects included in all specifications. Robust standard errors, clustered at the session level, reported. ***,**,* indicates significance at the 1%, 5% and 10% levels, respectively.

We observe a strong positive relationship between donation decisions and empathy. A higher IRI score is related to a lower likelihood of choosing not to donate and a higher

likelihood of finally donating. In contrast, the impulsivity measures are not significantly related to donation decisions. This suggests that, within the context of donation decisions made without time pressure or cognitive load, impulsivity plays a limited role in giving decisions.

6 Conclusion

It is well known that when the flow of utility is out of synch with the flow of consumption, interesting things can come of it. Here we build a nearly trivial model of intertemporal altruism around a simple, intuitive, yet extremely powerful observation. Whether a commitment to give is irrevocable, like putting money into a Donor Advised Fund, or a promise that can be broken, like a pledge to give $100 to Public Radio, at the moment the person decides on their intention they feel some of the joy of giving to come from the full giving experience. If the gift is made at the same time as the decision is made, there is no consequence of our assumptions. But if giving is to transpire in the future, then this assumption can mean that more dollars will be committed, more pledges will be made, and potentially more money will collected by charities who allow the collection of dollars to be later than the collection of intentions of the donors. All of these predictions rely only on well-understood and easily manipulated models of discounted utility.

We test our model and predictions with a laboratory experiment on giving to a real charity. In a study that lasts two weeks, subjects are asked in the first week if they would like to give $5 from their experimental earning to a charity called Give Directly. They are given a 15 minute presentation about the charity, including evidence of its effectiveness, and then are shown the photo of a household in Kenya that they could be helping with their donations. They are then given several minutes to make their giving decision.

In the Give-Now condition, subjects pay $5 from their first week's payment. In the Give-Later condition they make the same binding decision, but the $5 comes from their second week's payments. We find theoretically and experimentally that giving should rise by delaying the payment.

The third condition is called Pledge. The subject expresses that, yes, they will donate $5 to GiveDirectly in the second week of the study, and will confirm that decision "next week." Both theory and the experiment show that the number pledging should exceed the number committing to give in Give-Later, but that the final number of donations should be about the same or only slightly higher than in Give-Now.

The final condition is Pledge-or-Give-Now. This offers subjects either commitment by giving in the first week of the study, or flexibility by pledging to give in the second week of the study. Both the theory and experiment show that the number of people saying no to giving in the first week will be the same in both conditions, but many people will in fact choose commitment. According to our model, neither giving immediately nor pledging to give later are signs of dynamic inconsistency or its remedies, but are signs of individuals with different preference profiles choosing the timing of giving that suits them best.

Since our key assumption is in stark contrast to models of self-control and altruism, we also elicited subjects' attitudes about their experience in the experiment, as well as administered psychological scales for empathy and impulsivity. While these measures are noisy, there were clear aspects that supported our behavioral assumptions, while assumptions of impulsivity that would be necessary for a self-control problem were not predictive.

In sum, the model presented here is extraordinarily simple, and we see this as a strength. Not only are its assumptions natural, intuitive, and appealing, the model is easy to understand and to manipulate for generating predictions. Most importantly, the model's assumptions and its predictions are both validated by experimental data.

References

[1] Andreoni, J. (2006). Philanthropy. *Handbook of Giving, Reciprocity and Altruism*, L.-A. Gerard-Varet, Serge-Christophe Kolm and Jean Mercier Ythier, Editors, Elsivier/North-Holland.

[2] Andreoni, J., and Bernheim, D. B. (2009). "Social image and the 5050 norm: A theoretical and experimental analysis of audience effects." *Econometrica* 77, no. 5: 1607-1636.

[3] Andreoni, J., and Payne, A.A. (2013) Charitable Giving. *Handbook of Public Economics, Volume 5,* (Amsterdam: North Holland), Alan Auerbach, Raj Chetty, Martin Feldstein, Emmanuel Saez, editors, 1–50.

[4] Andreoni, J., and Rao, J.M. (2011). "The power of asking: How communication affects selfishness, empathy, and altruism." *Journal of Public Economics* 95, no. 7: 513-520.

[5] Andreoni, J., J. M. Rao, and H. Trachtman (2016). Avoiding the ask: a field experiment on altruism, empathy, and charitable giving. *Journal of Political Economy*, forthcoming.

[6] Bénabou, Roland, and Jean Tirole (2005). "Incentives and prosocial behavior."' *American Economic Review* 96, no. 5: 1652-1678.

[7] Barratt, E. S. (1959). "Anxiety and impulsiveness related to psychomotor efficiency." Perceptual and Motor Skills, 9, 191–198.

[8] Bekkers, R., and Wiepking, P. (2010). "A literature review of empirical studies of philanthropy: Eight mechanisms that drive charitable giving." *Nonprofit and Voluntary Sector Quarterly.*

[9] Breman, A. (2011). "Give more tomorrow: Two field experiments on altruism and intertemporal choice." *Journal of Public Economics*, 95(11), 1349–1357.

[10] Charness, G., and Dufwenberg, M. (2006). "Promises and partnership." *Econometrica*, 74(6), 1579–1601.

[11] Damgaard, M. T., and Gravert, C. (2014). "Now or never! The effect of deadlines on charitable giving: Evidence from a natural field experiment." Working paper.

[12] Dana, J., Cain, D. M., and Dawes, R. M. (2006). "What you don't know won't hurt me: Costly (but quiet) exit in dictator games." *Organizational Behavior and Human Decision Processes*, 100(2), 193-201.

[13] Dana, J., Weber, R. and Kuang, J.X. (2007). "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness." *Economic Theory* 33 (1), 67–80.

[14] Davis, M. H. (1980). "A multidimensional approach to individual differences in empathy." *JSAS Catalog of Selected Documents in Psychology*, 10, 85.

[15] Davis, M. H. (1983). "Measuring individual differences in empathy: Evidence for a multidimensional approach." *Journal of Personality and Social Psychology*, 44, 113–126.

[16] DellaVigna, S., List, J. A., and Malmendier, U. (2012). "Testing for altruism and social pressure in charitable giving." *The Quarterly Journal of Economics*, 127(1), 1–56.

[17] Dreber, A., Fudenberg, D., Levine, D. K., and Rand, D. G. (2014). "Altruism and Self-Control." Working paper.

[18] Ellingsen, T., and Johannesson, M. (2004). "Promises, Threats and Fairness." *The Economic Journal*, 114 (495), 397–420.

[19] Fiedler, S., Glöckner, A., Nicklish,A., and Dickhert, S. (2013). "Social Value Orientation and information search in social dilemmas: An eye-tracking analysis." *Organizational Behavior and Human Decision Processes*, 120, 272–284.

[20] Frederick, S. (2005). "Cognitive reflection and decision making." *Journal of Economic Perspectives*, 25–42.

[21] Gneezy, U. (2005). "Deception: The role of consequences." *American Economic Review*, 384-394.

[22] Huck, S., and Rasul, I. (2011). "Matched fundraising: Evidence from a natural field experiment." *Journal of Public Economics*, 95(5), 351–362.

[23] Knowles, S., and Servatka, M. (2015). "Transaction costs, the opportunity cost of time and procrastination in charitable giving." *Journal of Public Economics*, 125, 54–63.

[24] Kovarik, J. (2009). "Giving it now or later: Altruism and discounting." *Economics Letters*, 102(3), 152-154.

[25] Moeller, F. G., Barratt, E. S., Dougherty, D. M., Schmitz, J. M., and Swann, A. C. (2001). "Psychiatric aspects of impulsivity." *American Journal of Psychiatry*, 158, 1783–1793.

[26] Piovesan M. and Wengström, E. (2009). "Fast or fair? A study of response times." *Economic Letters*, 105, 193–196.

[27] Rand, D. G., Greene, J. D., and Nowak, M. A. (2012). "Spontaneous giving and calculated greed." *Nature*, 489(7416), 427–430.

[28] Rand, D. G., and Kraft-Todd, G. T. (2014a). "Reflection does not undermine self-interested prosociality." *Frontiers in behavioral neuroscience*, 8.

[29] Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., and Greene, J. D. (2014b). "Social heuristics shape intuitive cooperation." *Nature communications*, 5.

[30] Recalde, M.P., Riedl, A., and Vesterlund, L. (2014). "Error Prone Inference from Response Time: The Case of Intuitive Generosity." CESIfo Working Paper No. 4987.

[31] Serra-Garcia, M., Van Damme, E., and Potters, J. (2013). "Lying about what you know or about what you do?" *Journal of the European Economic Association*,11(5), 1204–122.

[32] Stanford, M.S., C.W. Mathias, D.M. Dougherty, S.L. Lake, N.E. Anderson, and J.H. Patton (2009). "Fifty years of the Barratt Impulsiveness Scale: An update and review." *Personality and Individual Differences* 47, 385–395.

[33] Tinghög, G., Andersson, D., Bonn, C., Böttiger, H., Josephson, C., Lundgren, G. and Johannesson, M. (2013). "Intuition and cooperation reconsidered." *Nature*, 498(7452), E1–E2.

[34] Trachtman, H., Steinkruger, A., Wood, M., Wooster, A., Andreoni, J., Murphy, J.J., and Rao, J.M. (2015). "Fair Weather Avoidance: Unpacking Costs and Benefits in a Replication of "Avoiding the Ask"." *Journal of the Economic Science Association*, Vol 1, 8–14.

Appendix A: Instructions

Note: the text in square brackets was not read aloud.

Welcome

Thank you for participating in this experiment. During the experiment you and the other participants are asked to answer a series of questions. Please do not communicate with other participants. If you have any questions please raise your hand and an experimenter will approach you and answer your question in private.

This experiment consists of two parts.

- Part 1: Today we will ask you to answer a series of questionnaires.
- Part 2: A follow up survey that you will be asked to fill out a week from today.

Payment

You receive for the participation in this experiment $26. Please note that in order to obtain you all payments you need to answer both parts of the experiment.

- Today you receive $6 for showing up to the experiment and answering the first part of the experiment. You can collect the $6 from the experimenter after the session is finished.
- The remaining $20 you will receive at the end of the next week?s session.

[GiveDirectly Pitch

Slides of GiveDirectly are shown on the screen. Experimenter reads the slides]

[At the end of the pitch:]

- [Condition Give-Now]: We would like to ask you whether you would like to donate $5 of your show up fee for today's session to GiveDirectly. You will be asked to answer this question on your screens in a minute. If you answer "YES, I'd like to donate $5 today," $5 of your show up fee today will be donated. If you say "NO," no donation will be made. Your decisions are final today.

- [Condition Give-Later]: We would like to ask you whether you would like to donate $5 of your show up fee for next week's session to GiveDirectly. You will be asked to answer this question on your screens in a minute. If you answer "YES, I'd like to donate $5 next week," $5 of your show up fee next week will be donated. If you say NO, no donation will be made. Your decisions are final today.

- [Condition Pledge]: We would like to ask you whether you would like to donate $5 of your show up fee for next week's session to GiveDirectly. You will be asked to answer this question on your screens in a minute. If you answer "YES, I'd like to donate $5 next week," we will ask you again next week and you can make your decision final at that time, then $5 of your show up fee next week will be donated. If you say "NO," no donation will be made, and that decision will be final today.

- [Condition Pledge-or-Give-Now]: We would like to ask you whether you would like to donate $5 of your show up fee to GiveDirectly. You will be asked to answer this question on your screens in a minute. If you answer "YES, I'd like to donate $5 today," $5 of your show up fee today will be donated. This decision will be final. If you answer "YES, I'd like to donate $5 next week," we will ask you again next week and you can make your decision final at that time, then $5 of your show up fee next week will be donated. If you say "NO," no donation will be made, and that decision will be final today.

Appendix B: Additional results

B.1. Analysis of attitudes toward the charity and happiness

At the end of the experiment, we elicited subjects' attitudes towards the charity as well as their perceptions about their decisions in the experiment. In a series of questions that were answered on a 5-item Likert scale from "strongly disagree" (1) to "strongly agree" (5), we elicited whether subjects liked GiveDirectly (*"I like the work of GiveDirectly"*); whether they planned to donate in the future to GiveDirectly (*"I plan to donate to GiveDirectly in the future"*); whether they regretted their donation decision (*"I regret my donation decision"*), or felt pressured to donate (*"I felt pressured to donate"*). We also elicited their happiness in general, regarding their participation in the study, about the opportunity to donate, and about their donation decisions. We examine in detail the latter two dimensions of happiness as these are the most closely related to the decisions in the experiment.[20] Additionally, we include in the analysis the willingness to sign up to an e-mail newsletter from the charity. This decision is costly to the recipient, if he actually does not wish the newsletter, and can therefore be contrasted with self-reported attitude measures.

We start by examining whether the different treatments had spillover effects on attitudes. Table B.1.1 displays the regression results when examining the impact of each treatment on attitudes. For five out of six attitude measures, as well as for the newsletter signup, we observe no spillover effects of the treatments. The only effect observed from delaying the donation (with and without commitment) is on happiness with the opportunity to donate (column (5) to Table B.1.1. The results suggest that individuals in these treatments are happier with the opportunity to donate. A possible interpretation of this effect is that in the treatments where donations are paid in $t = 1$, the decision to donate is more salient, as they just paid their donation, than in the Give Now treatment, where all donations occurred in $t = 0$.

Table B.1.2 displays the impact of thank-you notes on attitudes and newsletter signup.

[20] Happiness in general and with participation in the study are not significantly affected by the thank-you note.

Table B.1.1: Treatment effects on attitudes

	(1) Like charity	(2) Future donation	(3) Regret decision	(4) Pressure to donate	(5) Happy opportunity to donate	(6) Happy donation	(7) Newsletter
Give-Later	-0.100	-0.269	0.137	0.031	0.284**	0.070	-0.049
	(0.106)	(0.166)	(0.144)	(0.153)	(0.103)	(0.150)	(0.053)
Pledge	-0.228*	-0.255	0.183	-0.285	0.156*	-0.093	0.011
	(0.130)	(0.159)	(0.140)	(0.221)	(0.076)	(0.122)	(0.051)
Pledge-or-Give-Now	-0.130	-0.349**	0.117	0.041	0.259***	-0.031	0.025
	(0.115)	(0.153)	(0.101)	(0.141)	(0.072)	(0.109)	(0.052)
Constant	4.059***	3.401***	1.814***	3.083***	3.598***	3.637***	
	(0.166)	(0.267)	(0.121)	(0.220)	(0.116)	(0.123)	
Observations	458	459	458	458	458	457	453
R-squared	0.616	0.061	0.600	0.036	0.026	0.023	

Notes: this table presents the OLS regression results on the attitude questions regarding the charity and the donation decision (columns 1-6) and on marginal effects from probit regressions on sign-up to the charity newsletter (column 7). The questions in columns (1) to (6) were answered on a 5-item Likert scale from "strongly disagree" (1) to "strongly agree" (5). The variable Like charity measures the strength of agreement/disagreement with the statement "I like the work of GiveDirectly". The variable Future donationmeasures the strength of agreement/disagreement with the statement "I plan to donate to GiveDirectly in the future". The variable Regret decision measures the strength of agreement/disagreement with the statement "I regret my donation decision". The variable Pressure to donate measures the strength of agreement/disagreement with the statement "I felt pressured to donate". The variable Happy opportunity to donate measures the strength of agreement/disagreement with the statement "I feel happy about having the opportunity to donate". The variable Happy donation measures the strength of agreement/disagreement with the statement "I feel happy about my donation decision". Newsletter takes value 1 if the participant indicated that s/he wanted to be added to the charity's newsletter. The variables Give Later, Pledge, Pledge-or-Give-Now are dummies that take value one in the corresponding treatment, zero otherwise. All specifications include individual characteristics (asian, gender, native english speaker, cognitive reflection score) as defined in Table 1. Experimenter fixed effects included in all specifications. Robust standard errors, clustered at the session level, reported. ***,**,* indicates significance at the 1%, 5% and 10% levels, respectively.

The results reveal that the weak version of the thank-you note had no significant effects on any of the attitudes elicited, both in the Pledge and Pledge-or-Give-Now treatments. Similarly, the strong version of the thank-you note only led to a marginally significant increase in happiness with the donation. These results suggest that the strong version of the thank-you note did not have large effects on attitudes either. Hence, thank-you notes do not appear to have had strong spillover effects on attitudes. The effects of the thank-you note on newsletter signup are mixed. While we observe a negative effect of the weak thank-you note on signups in the Pledge treatment, this effect disappears in the Pledge-or-Give-Now treatment, where thank-you notes do not affect signups.

Table B.1.2: The impact of thank-you notes on attitudes

	(1) Like charity	(2) Future donation	(3) Regret decision	(4) Pressure to donate	(5) Happy opportunity to donate	(6) Happy donation	(7) Newsletter
Panel A: Pledge Treatment							
Thank-You: Weak	-0.140	0.028	-0.039	0.145	0.010	-0.117	-0.207**
	(0.091)	(0.184)	(0.167)	(0.139)	(0.157)	(0.251)	(0.086)
Constant	3.885***	3.103***	1.887***	2.348***	3.961***	3.976***	
	(0.108)	(0.347)	(0.209)	(0.272)	(0.337)	(0.243)	
Observations	72	71	71	71	71	71	72
R-squared	0.613	0.099	0.409	0.058	0.022	0.054	
Panel B: Pledge-or-Give-Now Treatment							
Thank-You:Weak	0.038	-0.003	-0.117	0.026	0.158	0.282	0.162
	(0.141)	(0.225)	(0.182)	(0.187)	(0.207)	(0.210)	(0.109)
Thank-You:Strong	-0.029	-0.168	0.268	0.383	-0.024	0.271*	0.081
	(0.149)	(0.230)	(0.180)	(0.345)	(0.202)	(0.136)	(0.073)
Constant	3.752***	2.945***	2.240***	3.033***	3.430***	2.955***	
	(0.272)	(0.282)	(0.360)	(0.249)	(0.239)	(0.170)	
Observations	95	95	95	95	95	95	73
R-squared	0.475	0.031	0.635	0.107	0.066	0.085	

Notes: this table presents the OLS regression results on the attitude questions regarding the charity and the donation decision, among those individuals who pledged to give, by email condition. The questions were answered on a 5-item Likert scale from "strongly disagree" (1) to "strongly agree" (5). The variable Like charity measures the strength of agreement/disagreement with the statement "I like the work of GiveDirectly". The variable Future donationmeasures the strength of agreement/disagreement with the statement "I plan to donate to GiveDirectly in the future". The variable Regret decision measures the strength of agreement/disagreement with the statement "I regret my donation decision". The variable Pressure to donate measures the strength of agreement/disagreement with the statement "I felt pressured to donate". The variable Happy opportunity to donate measures the strength of agreement/disagreement with the statement "I feel happy about having the opportunity to donate". The variable Happy donation measures the strength of agreement/disagreement with the statement "I feel happy about my donation decision". Newsletter takes value 1 if the participant indicated that s/he wanted to be added to the charity's newsletter. Since there is no variation in newsletter signup for one experimenter, in Panel B column (7), 22 observations are dropped. The variables Thank-You:Weak and Thank-You: Strong are dummies that take value one in the corresponding condition, zero otherwise. All specifications include individual characteristics (asian, gender, native english speaker, cognitive reflection score) as defined in Table 1. Experimenter fixed effects included in all specifications. Robust standard errors, clustered at the session level, reported. ***,**,* indicates significance at the 1%, 5% and 10% levels, respectively.