Time-Inconsistent Charitable Giving[*]

James Andreoni
University of California, San Diego and NBER

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

October 31, 2016

Abstract

This paper examines the interaction between moral contradictions and time in charitable giving. Applying a simple theoretical framework to two longitudinal experiments with actual charitable donations, we show that moral contradictions become the source of a new kind of time inconsistency linked to a demand for flexibility, rather than the more typical demand for commitment. This kind of time inconsistency coexists with the opposite of kind of time inconsistency arising from temptation to give, which is exhibited by a substantial minority of individuals. Our results reveal that time inconsistency is pervasive and exhibits unique features in the charitable domain.

JEL classification: D64, D90, C91.

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

1 Introduction

For many years research on charitable giving has focused on understanding how utility flows from the act of giving. We have learned that giving can be intrinsically joyful (e.g., Andreoni 1989, 1990; Ribar and Wilhelm, 2002), enhance self-image (e.g., Benabou and Tirole, 2006), improve social-image (e.g., Andreoni and Bernheim, 2009), and be a source of prestige (Harbaugh, 1998). Giving can also have less clearly positive effects. It can create moral contradictions for givers (e.g., Dana, Cain and Dawes, 2006; Dana, Weber and Kuang, 2007), and can be manipulated by social pressure and empathy. Some individuals enjoy saying yes to a fundraiser (Crumpler and Grossman, 2008; Andreoni and Rao, 2011; Chowdhury and Jeon, 2015), some dislike saying no, others will have a good or bad feeling just from being asked (DellaVigna, List, Malmendier, 2015; Trachtman et al., 2015; Andreoni, Rao, and Trachtman, 2016).[1]

These avenues for utility have been studied from a static perspective thus far, where act of giving occurred at a single point in time. In this paper we show that adding time to the decision to give enhances the importance of moral issues in giving. This has three main implications. First, if individuals do not enjoy giving, but do dislike saying no to a fundraiser, then introducing time in the giving process can generate additional giving: when a commitment to a future gift is solicited, individuals who commit avoid saying no immediately, while delaying the cost of giving. Second, pledges to give in the future are often made, but do not increase giving, since most pledges are insincere. Third, and most importantly, moral contradictions become a new source of dynamically inconsistent behavior, with features that are qualitatively different from dynamic inconsistencies that have been well documented in standard consumption choices (for a review, see, Frederick, Loewenstein and O'Donoghue, 2002). Dynamic inconsistencies that arise with moral contradictions are tied to the desire for flexibility and, hence, unlike other consumption choices, with a strict preference *not* to commit to choices made in advance.

In this paper, we conduct two longitudinal experiments to examine the impact of the introduction of time in actual donation decisions. To the best of our knowledge, these are the first

---

[1] For reviews, see Andreoni (2006), List (2011), and Andreoni and Payne (2013).

experiments to provide a comprehensive study of the dynamics of charitable giving. We also develop a stylized theoretical framework that captures the interaction between moral contradictions and time in the giving process. We find empirical support for the three implications above, all of which can be explained within our theoretical framework.

The theoretical framework we propose is based on the idea that charitable giving is a collection of decisions and transactions, each of which brings their own flows of utility, money and production of charitable goods. At the beginning of the process, after the ask has been made, the individual faces the decision whether to say yes or no to the fundraiser. Several recent papers suggest that individuals may derive utility from saying yes (Crumpler and Grossman, 2008; Andreoni and Rao, 2011; Chowdhury and Jeon, 2015) and disutility when saying no (DellaVigna, List, Malmendier, 2015; Trachtman et al., 2015; Andreoni, Rao, and Trachtman, 2016), such that some may be willing to incur costs to avoid the ask. The model thus assumes that individuals derive utility from making a giving decision, which is separate from the utility experienced when paying for the gift. Hence, we depart from standard models of giving, which typically assume that the utility from giving is only generated by the act of giving (e.g., Andreoni, 1989, 1990).

This simple change yields several predictions. We consider first giving choices when gifts are delayed, and individuals commit in advance to their giving decision. The first result is that more individuals will give when the decision to give is made in advance, than when the decision and payment of the gift both take place immediately. The increase in giving is generated by individuals who dislike giving, that is, those for whom the utility generated when making the gift is lower than its cost. They choose not to give when the gift is paid immediately, but choose to give when the cost of giving is delayed, as giving is made more attractive due to simple discounting of its cost, relative to the immediate cost of saying no. Second, these dynamically inconsistent individuals exhibit a strict preference for flexibility. If there is a future opportunity to revise their commitment, individuals choose to give in advance, but choose not to commit to that choice, leaving for themselves room to revise their choice towards saying no later.

We next consider giving choices when individuals pledge to give in the future, where pledges are non-binding. Pledges offer individuals an opportunity to delay the cost of say-

ing no. However, reneging may carry moral costs derived from breaking a promise or stated intention (e.g., Charness and Dufwenberg, 2006; Serra-Garcia, van Damme and Potters, 2013). This leads to three predictions. First, many will choose to pledge, but pledging will not increase ultimate giving significantly, unless the moral costs of reneging on a pledge are high. Second, offering both the opportunity to pledge and to give immediately are offered, will induce selection. Those types who choose to pledge are overwhelmingly those who are delaying the no, and hence will renege frequently. Third, having obtained a pledge from individuals who intend to renege, one may increase the "social pressure" to confirm the pledge and donate by thanking individuals who pledge, and thus decrease reneging.[2]

In our first experiment we examine the effect of introducing time in the giving process, while in the second we provide within-subjects evidence of dynamic inconsistency and its relation with commitment demand. In the first experiment, over 690 undergraduate students participated in a longitudinal experiment, that took place over two weeks, in which they were asked make an actual donation of $5 to GiveDirectly, a charity providing cash transfers to individuals living in extreme poverty in Kenya and Uganda.[3] In the baseline treatment, donations were made immediately in the first week of the study. In this treatment, 31% of subjects choose to give. In contrast, when individuals were asked to commit in the first week to donate in the second week of the study, donation rates increased to 45%, an increase of nearly 50% in giving. This provides evidence that a significant fraction of individuals are time inconsistent or, more precisely, violate time stationarity (Halevy, 2015).[4]

When pledging is introduced as the only form of giving, this results in a high rate of pledging, of over 65%. However, over 45% of individuals who pledge in the first week choose to renege on their pledge in the second week of the study, leading to an ultimate level of giving that is indistinguishable from that achieved when requesting immediate gifts. If donors are

---

[2]Thank-you notes are a frequently-used strategy by charities to cultivate repeat donors. Here, we consider their effectiveness within a single fundraising campaign, to increase the likelihood of converting pledges into gifts.

[3]Details regarding the procedures of the experiment are provided below.

[4]Evidence of violations of time stationarity is also provided in two large-scale field experiments by Breman (2011). Among employees who are already donating to charity, she finds that requests to increase donations are significantly more effective when the increases are delayed, rather than immediate. These experiments however do not provide evidence of within-subject dynamic inconsistency and commitment demand, two features that are crucial in understanding the source of violations of time stationarity.

offered the choice to give immediately or pledge in the first week, we observe 21% choosing to give immediately. There is evidence of selection among self-selected pledgers, who display a significantly higher rate of reneging, of over 70%, relative to the treatment in which pledging is the only form of giving. Further, when individuals who pledge are thanked for their pledge the same day they pledge, and a week before they are asked to confirm their pledge, reneging decreases and ultimate giving increases. These findings are all consistent with a significant majority of our subjects exhibiting moral contradictions as described in our framework, where individuals display a desire to avoid saying no to the fundraiser and, at the same time, to avoid giving.

Our second experiment examines in detail the source of the violations of time stationarity observed in our first between-subjects experiment. According to our framework, dynamic inconsistency is a result of moral contradictions, which yields the unique prediction that individuals who are dynamically inconsistent – that is, who choose to donate in advance, but not immediately – will not demand commitment, but will, in fact, have a preference for flexibility. This contrasts with dynamic inconsistency that is generated by temptation or self-control problems (e.g., Laibson, 1997; Fudenberg and Levine, 2006; and Gul and Pesendorfer, 2007). Dreber et al. (2016) show that individuals may face self-control problems in the social domain, where the short-run self is assumed to be altruistic, while the long-run self is assumed to be selfish. Their model predicts the opposite of type of dynamic inconsistency than we documented in our first experiment. Individuals in their model are less likely to give when the decision is made in advance, compared to when giving is immediate, and exhibit a preference for commitment.[5] Our second experiment hence uses a within-subjects design that identifies both dynamic inconsistency and commitment demand.

Our results reveal that 37% of the subjects are dynamically inconsistent, making different donation decisions depending on whether these concern immediate or delayed gifts. Among these, a majority (62%) exhibit the type of dynamic inconsistency documented in our first

---

[5]The predictions of Dreber et al. (2016) have found support using dictator games (see also, Kovarik, 2009), and closely relate to experiments manipulating time pressure and cognitive load in social dilemmas (e.g., Rand, Greene and Novak, 2012; Rand et al., 2014a,b), although other evidence in the latter settings is mixed (e.g., Kessler, Kivimaki, Niederle, 2016, and Recalde, Riedl, & Vesterlund, 2015).

experiment: they choose to give when the choice is made in advance, while they reverse their decision, choosing not to give, when giving occurs immediately. Most of these subjects (65%) exhibit a strict preference for flexibility, i.e., they do not wish to commit to the giving decision made in advance. This reveals that the source of dynamic inconsistency for these subjects is not temptation, but can be explained by moral contradictions experienced in the giving process.

At the same time, we find a significant portion of subjects (38%) who exhibit the opposite type of dynamic inconsistency. These subjects choose not to give in advance, but give when the gift is immediate. More than half (52%) of these subjects exhibit a strict preference for commitment. This suggests that temptation to give is an important driver of dynamic inconsistency for some of our subjects, in support of the model by Dreber et al. (2016).

This heterogeneity among givers is of extreme interest. It shows that introducing time in the process of giving can uncover important dynamics, as well as substantial heterogeneity in the motivations for giving. It also suggests the value in creating fundraising strategies that allow donors to reveal their types, and for fundraisers to tailor solicitations to each type.[6]

In order to establish the first main set of facts resulting from the introduction of time to the decision to give, we begin this paper by presenting our primary between-subjects experiment in Sections 2 and 3. Section 4 provides a simple framework that captures our main findings, including illustrating the dynamic inconsistency. In Sections 5 and 6 we present our within-subjects experiment, where we test for individual heterogeneity in dynamic inconsistency and relate dynamic inconsistency to commitment demand. Section 7 discusses how our ideas can be used to deepen the understanding of the giving process.

2   Experiment 1: Putting Time Into Giving Money

Our Experiment 1 consisted of two sessions spread exactly one week apart, where the opportunity to donate $5 to charity was presented. We refer to the first and second sessions as the week 1 and week 2 sessions, respectively.

In all treatments, the week 1 session opens with a scripted slide show about the charity

---

[6]See also Andreoni et al. (2016) for an exercise of this type related to labor supply.

GiveDirectly (www.GiveDirectly.org). The presentation, which lasts about 15 minutes, discusses the work the charity does by giving direct cash grants to poor households in Kenya and other African nations. It also discusses the results of scientific evaluations of the program (Haushofer and Shapiro, 2016) showing very high returns on investment as well as endorsements from charity rating groups such as GiveWell. Importantly, the presentation also highlights that one of the co-founders and current officers of GiveDirectly is Professor Paul Niehaus of the Department of Economics at the University of California, San Diego, where the study was conducted. This, we expect, adds confidence to both our claims about the quality and efficacy of the charity and our (true) promises that the donations would indeed go to GiveDirectly. The presentation ends with an ask to give $5. The experimental treatments vary on when the final decision will be required, and when financial transactions will occur.

2.1 Experimental Treatments

The baseline treatment is called Give-Now. Here, all steps take place at once, with the ask, the decision to donate, and gift happening in week 1. In the second treatment, Give-Later, the ask and the decision to donate occur in week 1, while the gift is delayed to week 2.

In the third treatment, Pledge, the ask takes place in week 1, at which point individuals can choose to pledge a donation. If they do, they state an intention to give a week later, subject to confirmation then. The Pledge treatment, therefore, allows subjects who intend to say no an opportunity to postpone announcing this decision.[7] The introduction of the delay naturally raises the question, whom does it benefit? Someone who is certain to give in the future might wish for a more clear signal of her intentions, while someone who is sure to say no in the future can use the delay to discount the uncomfortable feeling of saying no. Acknowledging this heterogeneity, we added a treatment that allows individuals to self-select into the timing they preferred. In the Pledge-or-Give-Now treatment, individuals could pledge in week 1 to

---

[7]The wording of the potential answers to the ask was always either "No," or "Yes, I'd like to donate $5 X," where for Give-Now X was "today," for Give-Later X was "next week," and for Pledge is was "next week. Ask me again next week and I will make my final decision." To avoid confusion about the level of commitment to the pledge, the word pledge was not used anywhere in Pledge treatment. However, the initial statement of an intention to give may have been viewed by subjects as a promise (see Hanfling, 2008, for a philosophical argument), which is our intention. The instructions are presented in Online Appendix A.

give in week 2, or decide to give immediately, in week 1.

The second question raised by pledging is, can the extra time between the pledge and the final decision can be used productively by the charity? Since fundraisers extoll the benefits of a carefully designed thank-you letter,[8] we hypothesized that one way to increase the conversion of pledges into donations may be to use gratitude. To explore this, in the Pledge and Pledge-or-Give-Now treatments we sent thank-you notes via email to a randomly chosen subset of subjects who pledged to give in the first week of the experiment. The e-mail was delivered by 5:00 p.m. on the same day of the session in week 1, seven days prior to having to confirm their pledges.[9]

To examine how the thank-you letter may be working, we designed both a "strong" and a "weak" version of the thank-you note. The weak thank-you note emphasized the importance of the pledge and thanked individuals for pledging. The strong thank-you note included two manipulations shown elsewhere to enhance the identifiable victim effect and to strengthen identity as a donor.[10] We compare the effect of the weak versus strong thank-you note in the Pledge-or-Give-Now treatment. In the Pledge treatment only weak thank-you notes were sent.

2.2 Procedures

A total of 692 students participated in Experiment 1, conducted at the UC San Diego Economics Laboratory. There were 180 students in the Give-Now treatment, 179 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 treatment to have enough observations when

---

[8]See, for example, "Tips for thanking (and keeping) donors," in the *Chronicle of Philanthropy*, December 22, 2015.

[9]All subjects received an email 24 hours prior to their week 2 session simply reminding them to attend.

[10]Specifically, 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 the recipient family. The note closed by stating that we looked forward to seeing them in a week when they could confirm their pledge. 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. This reflects the importance of the identifiable victim, as shown by Small & Loewenstein (2003). In addition, the weak note thanked them for their pledge, while the strong note thanked them for "being a donor," which is a framing device known by psychologists such as Bryan, Adams & Monin (2013) and Walton & Banaji (2004), to increase the appeal to an individual's identity as a donor and thereby increase behavior in line with this identity.

examining the effect of the thank-you note on giving.

A concern in a longitudinal experiment is attrition. For this reason, the first set of sessions had a higher show-up fee in week 2 than in week 1 ($6 in week 1 and $20 in week 2). The attrition rate was low, 8%, and did not vary with the treatment, the decision subjects made in week 1, or their individual characteristics. A detailed analysis of attrition is shown in Online Appendix B.

A drawback of the differential show-up fees is that they could affect giving decisions and contaminate our treatment effects, especially in Give-Now and Give-Later. Thus, we self-replicated Give-Now and Give-Later treatments with equal show up fees ($15 each week). We found no significant differences in the sessions with equal show up fees and those weighted toward week 2.[11] Given this, in what follows we pool the data, excluding those who did not participate in both sessions.[12]

3   Results For Experiment 1

Our main objective for Experiment 1 is to examine the treatment effects on ultimate giving. Subordinate to this, we also hope to see whether the treatments can reveal information about preferences in week 1 that can be useful to the charity in the time before week 2 decisions.

3.1   Week 1 Decisions

Figure 1 presents our results for week 1. The first two bars show that introducing a delay between the decision and the gift increases giving. In the Give-Now treatment, when deciding and giving occur together and immediately, 30.9% of the subjects choose to donate. In the Give-Later treatment, where only the payment of the gift is delayed, the giving rate rises

---

[11]First, attrition was 10.5% in the second set of sessions and not significantly different by show-up fee ($\chi^2 = 0.242$, $p = 0.623$). Second, donation rates were 32.5% and 29.4% in Give-Now ($\chi^2 = 0.184$, $p = 0.668$), and 43.8% and 46.5% in Give-Later ($\chi^2 = 0.114$, $p = 0.736$), in the first and second set of sessions, respectively. Third, individuals who did not participate in the week 2 session of the experiment did not behave differently in week 1 than individuals who did. Their donation rates in week 1 are not significantly different ($\chi^2 = 0.369$, $p = 0.544$). Detailed results for all participants are shown in Online Appendix B.

[12]Online Appendix B also provides detailed results including all participants, and shows the conclusions from the analysis remain unaltered.

significantly to 45.3% of the subjects ($\chi^2 = 7.104$, $p < 0.01$). This nearly 50% increase in the rate of giving resulted from a delay of just one week in payment. The data reveal a significant violation of time stationarity in giving choices made by subjects across the Give-Now and Give-Later treatments. Moreover, if this pattern remains in the within-subjects setting of Experiment 2, it will point to dynamically inconsistent choices: when choosing consumption for next week, the decision is different than the one that would have been made in a week for consumption immediately.

It is important to note the subtle but real difference here with other consumption choices. Receiving an object in the present but paying for it later is obviously better. But in our experiment, no actual goods or dollars were exchanged for another week–only the decision is made today.

Note: Error bars denote $\pm$ 1 S.E.

Figure 1: Giving and Pledging in Week 1

9

The final two bars of Figure 1 show the results of the pledging treatments. In the Pledge treatment, 65.5% of the subjects pledge to donate. In the Pledge-or-Give-Now treatment, the total percentage of subjects who pledge or give immediately is 69.5%. These frequencies are not significantly different from each other ($\chi^2 = 0.543$, $p = 0.461$), but are significantly higher than the frequency of giving in Give-Now ($\chi^2 = 31.860$, $p < 0.01$; $\chi^2 = 53.690$, $p < 0.01$, respectively), and Give-Later ($\chi^2 = 10.63$, $p < 0.001$; $\chi^2 = 21.36$, $p < 0.001$, respectively). While not significantly higher than the Pledge treatment, the Pledge-or-Give-Now shows the greatest participation, with 48.2% of the subjects pledging to donate and 21.3% of the subjects giving immediately.

3.2 Week 2 Decisions

Individuals who pledged in week 1 were asked in week 2 to confirm their donations or to renege. Figure 2 graphs the frequency with which individuals renege on their pledges.[13]

Consider first the case with no thank-you notes. Here, 46.9% of the individuals renege in the Pledge treatment. This fraction increases to 70.8% in the Pledge-or-Give-Now treatment ($\chi^2 = 3.214$, $p = 0.073$), confirming that the option to give immediately in the Pledge-or-Give-Now treatment induces selection. In particular, those least likely to renege appear to have selected giving in week 1, thus are not among those pledging to give, leading to a higher rate of reneging.

What happens when a thank-you note follows a pledge? In the Pledge treatment, individuals receiving a weak version of the thank-you note renege in 42.5% of the cases, compared to 46.9% when they do not receive a thank-you note. This drop in reneging is, however, small and not significant ($\chi^2 = 0.138$, $p = 0.710$). In the Pledge-or-Give-Now treatment, pooling all conditions with thank-you notes and comparing rates of reneging with and without thank-you notes, we find that the thank-you note reduces reneging by a very large and significant 22.9 percentage points ($\chi^2 = 3.798$, $p = 0.051$). When comparing the strong and the weak

---

[13]In the Pledge treatment 72 subjects pledged to give in week 2. Among pledgers, approximately half (55.6%) received the weak version of the thank-you note. In the Pledge-or-Give-Now treatment 95 subjects pledged to give in week 2. Among them, 27.4% received the weak version of the thank-you note and 47.4% received the strong version.

Note: Error bars denote ± 1 S.E.

Figure 2: Reneging in week 2

versions of the thank-you notes, however, we find the strong thank-you performed slightly better on average, with reneging about 5 percentage points below that for the subjects getting the weak thank-you, but the difference in not significant ($\chi^2 = 0.734$, $p = 0.786$). Thus, while the identifiable victim and the identity manipulations may be effective on their own, it would appear that their effects are largely overwritten when presented alongside a message of gratitude.

The stronger effect of the thank-you note in Pledge-or-Give-Now further confirms the role of self-selection in this treatment. It suggests that many of the individuals who pledged appear to have done so with the intention of reneging, and for them the thank-you note resulted in a significant reduction in reneging. This effect is especially striking in light of the fact that the thank-you note came within a few hours of their pledges and a full seven days before subjects returned to confirm them or renege.

Table 1: Ultimate giving by treatment

Treatment	Donation rate	Std. Error
Give-Now	0.309	0.036
Give-Later	0.453	0.039
Pledge	0.364	0.046
Pledge+Without thank-you	0.348	0.068
Pledge+Thank-you:Weak	0.376	0.062
Pledge-or-Give-Now	0.437	0.035
Pledge-or-Give-Now+Without thank-you	0.354	0.068
Pledge-or-Give-Now+Thank-you: Weak	0.454	0.068
Pledge-or-Give-Now+Thank-you: Strong	0.470	0.052

3.3 Ultimate giving

Table 1 presents the rate of giving by treatment, and Table 2 presents the results of the regression analysis of the treatment effects.[14] Column (1) of Table 2 shows the results of a probit regression, ignoring the presence of thank-you notes. To examine the effect of thank-you notes, we conduct a placebo test by first assigning those who did not pledge to a thank-you condition with a probability equal to that of their counterparts who did pledge. We then examine the effect of the thank-you conditions using a weighted probit regression. This analysis is shown in column (2) of Table 2.[15]

The regressions in Table 2 reconfirm that separating the timing of giving from the decision to give (Give-Later) significantly raises the donation rate, leading to a 14 percentage point increase ($p \leq 0.001$), when a one week delay in the gift is introduced. This predicts an increase in donations of almost 50% relative to Give-Now.

Table 2 also shows that pledging alone has a small effect on giving, of less than 5 percentage points, that is not statistically significant. For pledging to show a measurable positive and significant effect on giving in our experiment, it must be accompanied by both the option to give now, in week 1, and a thank-you note of either kind. When these conditions are both met,

---

[14]Our analysis of the treatment effects in Table 2 reports p-values that are uncorrected for multiple hypothesis testing (e.g., List et al., 2016). However, since all p-values for significant differences are below 0.001, correcting p-values leaves our conclusions unchanged.

[15]An alternative approach is to randomly assign a share of the individuals who did not pledge to each thank-you condition, and use bootstrapping. Results remain qualitatively similar with this approach.

Table 2: Treatment effects on ultimate giving

	(1)	(2)
	Donation	
Give-Later	0.145***	0.144***
	(0.041)	(0.041)
Pledge	0.057	
	(0.049)	
Pledge+Without thank-you		0.042
		(0.056)
Pledge+Thank-you: weak		0.069
		(0.066)
Pledge-or-Give-Now	0.129***	
	(0.028)	
Pledge-or-Give-Now+Without thank-you		0.050
		(0.056)
Pledge-or-Give-Now+Thank-you: weak		0.147***
		(0.036)
Pledge-or-Give-Now+Thank-you: strong		0.161***
		(0.033)
Observations	631	631
Pseudo R-squared	0.011	0.0133

Note: This table presents the average marginal effects (calculated at the means of all variables) from probit regressions on ultimate giving decisions. Column (1) presents the marginal effect from simple probit regressions on the treatment. Column (2) presents results from weighted probit regressions, whereby individuals who did not pledge in Pledge and Pledge-or-Give-Now are assigned to both the no thank-you and the thank-you conditions, and weighted correspondingly. The variables Give-Later, Pledge, Pledge+Without thank-you, Pledge+Thank-you, Pledge-or-Give-Now, Pledge-or-Give-Now+Without thank-you, Pledge-or-Give-Now+Thank-you are dummies that take value one in the corresponding treatment or treatment+thank you condition, zero otherwise. Robust standard errors, clustered at the session level, were used in each individual regression. ***,**,* indicates significance at the 1%, 5%, and 10% levels, respectively.

giving rises by 14.7 (weak) to 16.1 (strong) percentage points, a statistically and economically significant increase.

4 A Simple Framework for Analysis

Experiment 1 yields six main findings: (1) Individuals are more likely to give when the gift is paid later, as in Give-Later versus Give-Now, which suggests dynamic inconsistency in giving; (2) Far more intentions to give are expressed in the treatments that allow pledges as compared to those that require immediate giving, as in Pledge versus Give-Later; (3) Reneging in the Pledge treatment is frequent, with final giving below Give-Later but above Give-Now; (4) The

fraction of subjects who either pledge or give now in Pledge-or-Give-Now equals the fraction of those who pledge in the Pledge treatment; (5) For subjects in the two pledging treatments, a larger fraction of those pledging in Pledge-or-Give-Now end up reneging in week 2 than those pledging in Pledge; (6) Thank-you notes sent immediately after a pledge significantly increase the likelihood that a pledge is converted to a donation in Pledge-or-Give-Now.

These six findings cannot be explained by models of giving that assume that the utility of giving all flows from a single event (such as Andreoni 1989, 1990). In this section we write down a simple and parsimonious model of the giving process that allows us to identity several core components of the full process of giving. The model, though simple, neatly accounts for all six findings from the experiment just presented. Having said that, we also hasten to add that we view the model as incomplete, leaving flexibility for future research to flesh out the foundations of the components we put forth as well as to discover new ones.

4.1 Give Now

Consider a situation in which someone has been asked to make a charitable gift of a particular size, to which she can say yes or no. Since giving is measured in dollars, normalize utility across individuals so that the utility cost of paying for the gift is the same across people. Call this utility value g.[16] We assume the act of giving to charity also brings some (perhaps altruistic) utility, which we will call α.[17] This is the most typical formulation in the literature.

Now we generalize this model slightly to include the social utility surrounding being asked to give. It is well established that people dislike saying no to reasonable requests to give,[18] even while they may appreciate being asked.[19] Assuming the charity is seen by all as providing some benefits to society, we would expect saying yes to the ask brings social benefits, while saying no brings social costs. Notice, however, that saying yes and no are mutually exclusive.

---

[16]This is equivalent to normalizing utility by the marginal utility of money, as is commonly done in the Public Economics literature.

[17]We do not rule out that $\alpha \leq 0$ for some, as in the case of a donor being unsure about the legitimacy of the solicitation, although in our study we generally regard $\alpha \geq 0$.

[18]See Dana, Cain, and Dawes (2006), Dana, Weber, and Kuang (2007), DellaVigna, List, and Malmendier (2012), Andreoni and Rao (2011), and Andreoni, Rao, and Trachtman (2016).

[19]See Crumpler and Grossman (2008), Andreoni and Rao (2011), and Chowdhury and Jeon (2014).

As such, the decision between the two will depend on the *difference* in their values rather than their absolute amounts. To make this point clear, suppose we separated the utility from saying yes, call it $s_y > 0$, from the disutility of saying no, $-s_n < 0$. Then a person will say yes if $Z_y + s_y > Z_n - s_n$, where the Z's represent other utility consequences of the decision. Since "yes" and "no" are mutually exclusive, normalize the utility of yes to be 0 and the utility of no then to be the net values of their utilities, $-n = -s_n - s_y < 0$. Then $-n$ will be sufficient for modeling both "yes" and "no" choices.

We begin with the baseline case in which the decision to give and the act of giving are made simultaneously and both occur in the present. Since the utility cost for paying for the donation is fixed at $-g$, saying yes yields a net utility of $\alpha - g$. Saying no yields $-n < 0$. Combining these, an individual says yes to the request to give if

$$\alpha - g \geq -n.$$

Rearranging this we can define a critical level of generosity $\alpha_N(n) \equiv g - n$, such that an individual will give now if $\alpha \geq \alpha_N(n)$. Note that, if asked, an individual may decide to give even if giving yields a disutility, $\alpha - g < 0$, as long as this disutility is smaller than that of saying no. This point will be important later.

4.2   Give Later

Now we take the model above and put time between the two moments of decision and payment. As above, the person is asked at time $t = 1$ (period 1) to commit to giving, but the payment for the gift comes in the future, $t = 2$ (period 2). With $\delta \leq 1$ as the discount factor, then an individual says yes in $t = 1$ if the discounted utility of giving later exceeds the utility of saying no:

$$\delta(\alpha - g) \geq -n.$$

Rearranging again yields the critical value $\alpha_L(n) \equiv g - n/\delta$ such that those with $\alpha \geq \alpha_L(n)$ will give later. Since $\delta < 1$, then $\alpha_L(n) < \alpha_N(n)$ for all n. That is, anyone who would give now will also give later, but many who would say no to giving now will instead be willing to say yes to giving later.

4.3   Dynamically-Inconsistent Giving

An interesting implication of our first two predictions is that they raise the possibility of dynamically-inconsistent giving (result 1). Begin with those who say yes when asked in period 1 to give in period 2. Suppose in period 2 we set aside their earlier choice and ask them if they are willing to give now. Assuming their tastes for giving now are the same in period 2 as in period 1, then our data suggests that a significant sum of those saying yes in Give-Later would change their preferences when asked again in period 2 if they would like to give now.

What kind of preferences would display such inconsistency? They would need to satisfy

$$\alpha_N(n) > \alpha > \alpha_L(n).$$

Substituting in the definitions of α_N and α_L and rearranging, we see

$$-n > \alpha - g > -n/\delta. \tag{1}$$

Equation (1) reveals that dynamically-inconsistent types are those who are made worse off by having been asked, that is $\alpha - g < 0$. As a result, their discounted utility, $\delta(\alpha - g)$ is higher than their undiscounted utility. Stated differently, while allowing people to give later may encourage more giving, it encourages it among those whose utility from giving is already negative.

Dynamic inconsistency arises because individuals face a moral contradiction in choosing whether to give or not: they do not desire to give, but also dislike saying no to the fundraiser. If these types are sophisticated, they will not demand commitment to their advance choices. The reason is that they anticipate their desire to say no to giving when asked again in period 2 and, hence, are better off by not demanding commitment. Thus, dynamic inconsistency is

directly linked to demand for flexibility in giving decisions, in contrast to the standard link between dynamic inconsistency and demand for commitment generated by quasi-hyperbolic preferences (such as in Laibson, 1997; O'Donoghue and Rabin, 1999) and temptation models (Fudenberg and Levine, 2006; Gul and Pesendorfer, 2007). We will examine this difference in detail in our second experiment.

4.4 Pledging

Next we put time between the ask and the final decision to give by allowing donors to submit non-binding pledges. Pledges are common in fundraising campaigns in the United States but have received little study by economists.[20] When giving occurs through pledges the decision to give can be split into two decisions separated by time. First is the decision to pledge or say no to the ask. Second, for those who pledge, there is the later decision to honor the pledge or renege.

Since pledging allows a non-giver to postpone saying no today, why should we not expect all non-givers to simply pledge and renege? Research shows that breaking pledges may entail costs akin to those from lying (see Ellingsen and Johannesson, 2004; Gneezy, 2005; Charness and Dufwenberg, 2006, and Serra-Garcia et al., 2013). Therefore, having pledged, breaking the pledge could have a utility loss of $\ell \geq 0$, in addition to the cost of saying no. Then in period 2 the full cost of reneging is $n + \ell$.

An individual who pledged will confirm the pledge if the utility from confirming is greater than from reneging: $\alpha - g \geq -n - \ell$. This leads to a critical $\alpha_C(n + \ell) \equiv g - (n + \ell)$, where those with $\alpha \geq \alpha_C(n + \ell)$ confirm the pledges.

When deciding to pledge, therefore, the potential donor compares three options. First is to pledge with the intention of giving. A person will choose this if $\alpha - g > -n - \ell$. Second is to

---

[20]Notable exceptions include Breman (2011), who finds an increase in giving from pledging. She revisits those who agreed to increase their monthly automatic payment to the charity a year after the pledge and found that virtually no one had reversed or reduced the increased monthly payment. Other studies of pledges are often more like our Give-Later treatment in that the are binding commitments to give, such as Zellermayer (1996), finds the result is sensitive to whether the pledge is framed as coming from currently held money (a negative effect), or from money earned in the future (positive effect). The psychologists Meyvis, Bennett and Oppenheimer (2010) find that those more capable abstract thinking are also more willing to give in the future.

pledge with the intention of reneging, implying $\alpha - g < -n - \ell$ and $-n < -\delta(n + \ell)$. Third is simply to say no immediately. This happens if $\alpha - g < -n - \ell$ and $-n > -\delta(n + \ell)$.

We can easily see from this analysis that the utility of pledging lies between the utility of Give-Now and Give-Later treatments, depending on the value of ℓ. If $\ell = 0$ then there is no penalty for pledging, and any subject will pledge. Those for whom $\alpha - g > -n$ will do so with the intention of giving, while the rest will have the intention of reneging. Thus, the situation reverts to that in Give-Now, the difference being that it is now played in period 2 rather than period 1. Suppose instead that reneging on a pledge is extremely costly, then any subject who pledges also gives. Specifically, if $\ell > n(1 - \delta)/\delta$, the environment mirrors that of Give-Later in which pledges are binding.

This now allows us two conclusions. First, the number of pledges in the Pledge treatment will be at least as high as the number of donors in Give-Later, and the lower the costs of lying, ℓ, the greater the number of insincere pledges. Second, the number of final donors in the Pledge treatment will be bounded below by the Give-Now treatment, and bounded above by Give-Later. The lower the costs ℓ, the closer final donations come to match the level in Give-Now. These predictions for pledging are consistent with our experimental data, results 2, and 3.

4.5   Pledge or Give Now

Recall, a person who donates in Give-Now satisfies $\alpha - g > -n$. Notice this inequality can be satisfied when both sides are negative as well as when the left side is positive. If we ask a donor in the Give-Now treatment to instead give later, her utility would become $\delta(\alpha - g)$. Thus only if she has $\alpha - g < 0$ would she agree; if $\alpha - g > 0$ she would prefer to give now.

Suppose that we offer subjects this choice in period 1: pledge to give next period or give now. The option to give now cannot make anyone worse off, and can only make them better off if they choose to give now. Thus, offering the choice to pledge or give now will not increase the total number of people willing to express an intention to give (by pledging or giving now) in period 1 above what we saw in the Pledge treatment, nor will it affect the ultimate number of donations collected. All it will affect is timing for those with high values for α, removing them from the pool of those pledging, and in doing so, will increase the average likelihood that a

pledge will be reneged. Again, all of these predictions are in line with our experimental results 4 and 5.

4.6   Gratitude

Seeing how significantly pledges reduce the number of people saying no immediately, it is natural to ask if the time between the pledge and the final giving decision can be used constructively to increase final giving. This leads to a tactic often used by fundraisers: showing gratitude.

Within our model, a thank-you note between the time of the pledge and the time for the donation could have two effects. First, the thank-you note could increase α. This would happen if gratitude increases the pledger's support or attachment to the organization. A second, but not mutually exclusive, effect of gratitude could be to increase n or ℓ. After being thanked a person may feel greater disutility from reneging, such as additional guilt from telling a lie. Either or both of these effects could lead to less reneging and are consistent with our sixth finding of our experiment: The thank-you note significantly increases the conversion pledges into actual donations.[21]

5   Experiment 2: Dynamic Inconsistency and Commitment

Perhaps the most interesting result of our study so far is the prediction of dynamic inconsistency in donations. In particular, more people agree now to give in the future than agree now to give now. If this effect were shown within subjects we would call this dynamically inconsistent behavior.

---

[21]The effect a thank-you note always differs across the Pledge and Pledge-or-Give-Now treatments due to selection effects. It is larger in the Pledge-or-Give-Now treatment because thank-you notes are sent to subjects with a significantly higher likelihood of reneging. Hence, a significant effect of thank-you notes will be detected more often in the Pledge-or-Give-Now treatment, than in the Pledge treatment, as is the case in our experiment. To illustrate this, suppose all subjects who gave in week 1 in the Pledge-or-Give-Now would have been forced to pledge and confirm in week 2 (like in the Pledge treatment). Then, the frequency of reneging among subjects who pledge and do not receive a thank-you note would have been 47.2%, while the frequency of reneging among those who receive a thank-you note would have been 32.4%. This difference in reneging is not statistically significant (χ^2-test, p-value=0.11), and much smaller in magnitude than that observed in the Pledge-or-Give-Now treatment.

Notice that the dynamic inconsistency generated in our theoretical framework does not rely on a non-exponential discounting model (such as in Laibson, 1997; O'Donoghue and Rabin, 1999), but rather stems from altering the timing of the social and economic payoffs of giving. This source of dynamic inconsistency is systematically different from inconsistencies resulting from temptation or present bias. A potential donor might be aware that his giving decision in advance is generated by picking the least of two bads, and be well aware that when giving is no longer delayed his decision would change. He would thus value flexibility over commitment.[22]

Dynamic inconsistency could also be of a different kind if in period 1 a subject chooses not to give in period 2, but when period 2 arrives she prefers to give. This version of dynamic inconsistency has been suggested by Dreber et al. (2016), and is based on the assumption that giving is tempting. In their paper subjects play dictator games with other subjects. The dictator makes decisions today about payments at different points in the future. They observe giving is highest when the gift is made in the present. This finding is consistent with giving being tempting.[23,24] While our between-subjects design does not point to a predominance of this effect, it also does not allow us to identify if a minority of subjects show dynamic inconsistency from temptation to give. Furthermore, our simple model does not consider this type of preference. Crafting the model to include temptation is trivial, but requires one extra degree of freedom.

---

[22]In a recent study Andreoni et al. (2016) find that people switch their decisions about how to allocate goods fairly between two people depending on the context of that decision. When the decision is ex ante to a partial realization of uncertainty, subjects favor ex ante notions of fairness. After the partial realization of uncertainty, subjects prefer to adopt an ex post stance on fairness, despite the fact that from the ex ante perspective this change is very unfair. Moreover, subjects rejected opportunities to commit to a fairness perspective, but rather preferred the flexibility for their fairness criterion to fit the context.

[23]This hypothesis relies on results found by Rand and coauthors (2012, 2014a, 2014b), who showed in various games that "fast decisions" also tended to be more altruistic. Recent work by Recalde, Riedl, and Vesterlund (2015), however, shows that errors are also correlated with time, and that earlier decisions are also more prone to errors, which undermines the inference that faster decisions are more representative of true preferences. Kessler, Kivimaki, and Niederle (2016), also use reaction times to infer base motives, and have mixed conclusions. That giving declines with time was shown by Meyvis, Bennett and Oppenheimer (2010) but attributed to a different source. They hypothesized that the present is more concrete while the future is more abstract. When the payment was framed as coming from currently held assets, pledges resulted in reduced giving, while opposite effect was found when the payment was framed as taken from future income. They found increased giving was associated with psychological measures of abstract reasoning.

[24]Note that an alternative model of temptation could be to assume that it is tempting to be selfish and, hence, that it is the long-run self who wishes to be altruistic. However, such a model would predict individuals do not avoid the ask, contrary to what has been documented in DellaVigna, List and Malmendier (2015) and Andreoni, Rao and Trachtman (2016). Such a model would also predict commitment demand, a finding for which we do not find support in Experiment 2.

For example, we could add a variable β_t, where $\beta_t > 0$ if the gift is transacted at the same time the decision is made, and $\beta_t = 0$ otherwise. With sufficient heterogeneity in the size of β—low temptation for some and high temptation for others (as was observed by Dreber et al., 2016)—then this approach would allow for the existence of two opposite types of dynamic inconsistency. How would a subject with such temptation preferences feel about commitment? Since temptation rather than social pressure is the motivation, the decision maker in these models has only her future self to answer to. Since this future self wishes the present self did not yield to temptation, commitment to the advance choice not to give is desirable.

In sum, depending on whether the donor feels she suffers from a bias that needs to be controlled or whether she is responding to changing social and economic incentives, different degrees of commitment demand for the two kinds of dynamically-inconsistent giving may arise. Next we describe our second experiment designed to answer all three of the questions raised: Is there dynamic inconsistency within subjects? If so, is the pattern (Donate, Not Donate) as the model of Section 4, is it (Not Donate, Donate) as suggested by a temptation to give, or do we see two different types of dynamic inconsistency? Finally, when there is dynamic inconsistency, is there also a demand for commitment, and does it differ for the two types?

5.1 Experimental Design

We designed a longitudinal experiment following the same structure and features as the Give-Now and Give-Later treatments. In week 1 subjects made a choice like that in Give-Later: a decision today about a gift paid in week 2. In week 2 they returned to the lab to make a second decision like that in Give-Now: a decision now about a gift paid now. In week 2, after both choices are made, the experimenter randomly selected either the Give-Later (period 1 decision) or the Give-Now (period 2 decision) to carry out.

We added one more step to the above in order to measure commitment demand. Knowing that either their week 1 or week 2 choice would be implemented at random, each subject was offered a probabilistic commitment device in week 1 (as in Augenblick, Niederle and Sprenger, 2015). In particular, after making her period 1 decision, the subject was offered the choice of

which week's decision would be *more likely* to be implemented. Let p be the probability week 1's choice is implemented. Then the subject chooses p from the set $p \in \{0.9, 0.1\}$. If the subject prefers week 1's outcome, she chooses the high p of 0.9, while if she prefers week 2's decision (which she has yet to make) she would choose the low p of 0.1. In order differentiate strict preferences for the earlier or later decision from a simple demand for randomization, we also allowed the subject to indicate indifference between high and low probabilities. If they did, their p was selected as either high or low with the flip of a coin, effectively creating a $p = 0.5$.

There were 183 participants in this experiment. Subjects were paid a show-up fee of \$15 in both week 1 and week 2 of the experiment. Twenty students (11%) did not participate in week 2. As before, attrition was unrelated to decisions made in week 1 ($\chi^2 = 0.750$, $p = 0.386$).

6   Results from Experiment 2

The share of subjects who decide to give in period 1 is 47.9%, a share that is similar to the 45.3% we observed in Give-Later (χ^2-test, $p = 0.644$). The share of individuals who decide to give in period 2 is 39.3%, a portion that is higher than the 30.9% in Give-Now, but not significantly different (χ^2-test, $p = 0.113$). This reveals dynamically inconsistent behavior is also present within subjects, though to a somewhat smaller extent than would have been predicted by the between subjects data, perhaps owing to concerns about appearing consistent (Falk and Zimmerman, 2016).

At the individual level, we find that 36.8% of the individuals are dynamically inconsistent by making different donation decisions in week 1 and week 2. Out of these, 62% choose (Donate, Not Donate) and 38% choose (Not Donate, Donate). The former type of dynamic inconsistency is significantly more frequent than the later (McNemar's test, $p = 0.07$). Thus, while both types exist, those apparently driven by social costs of saying no outnumber those who are hypothesized to face temptation to give in the present.

What about commitment demand? On aggregate, 35.4% of the individuals choose not to commit, 27.0% are indifferent, and 38.9% exhibit a strict preference for commitment. Figure

3 illustrates the choice of p in week 1 for each of the four possible dynamic paths. Here we see a clear difference between the two categories of dynamic inconsistency. Subject who choose (Donate, Not Donate) overwhelmingly prefer not to commit: 64.9% strictly favor flexibility while only 18.9% choose to commit. This shows that there exists a clear preference to maintain flexibility for giving to change in period 2. In contrast, among those who choose (Not Donate, Donate) we see far more people selecting commitment: 52.2% select commitment while 34.8% strictly prefer not to commit. This aligns with the prediction that commitment demand is used to overcome temptation to give.[25]

Note: The error bars denote $\pm$ 1 S.E.

Figure 3: Dynamic inconsistency and commitment demand

Looking at the two groups that were time consistent, we have no prediction for commitment for either group. Indeed, for (Donate, Donate) the commitment choices appear as if they were made at random, while (Not Donate, Not Donate) appears to favor commitment.[26]

---

[25] See Online Appendix B for detailed data that supports Figure 3.

[26] Dreber et al. (2016) predict that strict, rather than probabilistic, commitment can lead to commitment demand even by those who would otherwise resist donations in both periods. The reason is that the commitment saves the person the cognitive resources spent on exerting self-control. While in our experiment a decision is required in

Table 3 provides support for these results using a multinomial probit regression that relates commitment choice with the type of dynamic preferences expressed. As can be seen in column (1), individuals who choose to (Donate, Not Donate) are significantly more likely to choose no commitment than subjects in other treatments. Column (2), by contrast, shows that that those who choose (Not Donate, Donate) are marginally significantly more likely to commit than to express indifference. For dynamically consistent subjects, column (3) confirms that individuals who consistently donate over time do not exhibit a preference for flexibility or commitment, while column (4) reveals that individuals who consistently say no exhibit a marginally lower preference for flexibility, though they do not significantly prefer commitment, relative to indifference.

Table 3: Dynamic inconsistency and commitment demand

	(1)	(2)	(3)	(4)
	(Week 1 donation decision, Week 2 donation decision)			
	Dynamic inconsistency		Dynamic consistency	
	(Donate, not donate)	(Not donate, donate)	(Donate, donate)	(Not donate, not donate)
Flexibility: $p = 0.1$ for week 1 decision	0.259*** (0.098)	0.108 (0.103)	-0.110 (0.090)	-0.257* (0.138)
Commitment: $p = 0.9$ for week 1 decision	-0.040 (0.117)	0.141* (0.080)	-0.117 (0.097)	0.016 (0.154)
Observations		163		

Note: This table presents the marginal effects (calculated at the means of all variables) from a multinomial probit regression relating patterns of dynamic choice to commitment choice. Robust standard errors, clustered at the session level, are shown in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

The central finding of our second experiment is that there is heterogeneity surrounding dynamic inconsistency. Most dynamically inconsistent individuals are well described by the simple model we proposed, in which individuals agree to give later but not in the present. Since they may want the decisions they make to match the social payoffs at the time of the decision, they do not see these inconsistencies as problematic and thus do not choose commitment. A significant minority, however, also appear to adopt an opposite stance: They choose not to give

both periods regardless of commitment, one can imagine an interpretation their model that would cause the same effects here. If this is the case, then the fraction of subjects who satisfy their model would be under-estimated, and should include some time-consistent non-givers who nonetheless demand commitment.

in the future but cannot resist giving in the present. Those who are sophisticated about their self-control problem opt to commit to their period 1 choice to not give in period 2.

7 Discussion and Conclusion

Charitable giving can often be a complex social interaction with potentially many opportunities for social as well as economic payoffs to accrue. Seemingly minor changes in the timing or emphasis of the different components of the giving process can have significant effects on ultimate giving. A primary contribution of this paper is to demonstrate this fact about giving and point to the implications this has for dynamic behavior in the charitable domain. Most importantly, we document a new kind of dynamic inconsistency that is linked with demand for flexibility, in contrast to traditional sources of dynamic inconsistency which are linked with demand for commitment.

Understanding the socially complex process of charitable giving is necessary to answer many important questions about institutional design, the effect of tax policy on giving, and the welfare effects of both private and public efforts to affect giving. Our work illustrates the existence of flows of utility from different aspects of the decision, and hence suggests important ways in which these aspects can be manipulated to engineer outcomes. This finding should be of interest to fundraisers, economic theorists, and policy makers, as it suggests a new architecture for building models of donors, fundraising, and of tax policy toward giving. Consider the four examples below.

First, technology has made it easier to focus on "micro-donations" such as asking for the $1 gift at the checkout counter, or with purchases on Amazon.com. Dozens of companies have emerged in an attempts to disrupt the timing of saying yes and feeling the pain of paying, such as Google's OneToday app. A second example is the charitable sector's reaction to the concentration of wealth around the world. "High net worth donors" are often engaged in long term planning for legacy gifts, and are becoming the primary focus of fundraisers. These givers are deeply concerned with timing of the announcements of commitments to give, timing of the actual payments, longevity of naming rights, and other social and temporal aspects of the gift.

A third example is the rise of donor advised funds (DAFs) as a vehicle for de-linking the tax benefits of giving from the actual time of making the gift. DAFs act like reverse-401k funds for donations—the tax benefits of giving are realized when the money is set aside for charity, while the donations can be made at any future date, almost without restriction. Thus, taxpayers are loaning the tax benefit to the donor, at zero interest, until the money is, at some uncertain point, productively used by a charity. These funds present unique challenges for fundraisers and policy makers (see Andreoni, 2016).

Fourth, for the extremely wealthy, there is the "giving pledge." This pledge was started by Bill Gates and Warren Buffett in 2010, to encourage the world's wealthiest people "to give the majority of their net worth to philanthropy, either during their lifetime or upon their death" (givingpledge.org). This pledge challenges the traditional notion that the wealthy feel less rich by giving their money away, and raises the question of whether there is an asset value in past giving especially when such giving can be truly transformative. Stated differently, do the wealthy feel less rich by giving money away or are these gifts capitalized into their psychic notions of wealth?

In sum, more complete and nuanced models of giving, such as those pointed to here, will be needed to address the ever more complex world of charitable giving.

References

[1] Andreoni, J (1989). Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence. *Journal of Political Economy*, v. 97, December, 1447-58.

[2] Andreoni, J. (1990). Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *Economic Journal*, v. 100, June, 464-477.

[3] 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.

[4] Andreoni, J. (2016). Who Benefits from Donor Advised Funds. Unpublished working paper, 2016.

[5] Andreoni, J., Aydin, D., Barton, B., Bernheim, B. D., & Naecker, J. (2016). When Fair Isn't Fair: Sophisticated Time Inconsistency in Social Preferences. Available at SSRN 2763318.

[6] Andreoni, J., & Bernheim, D. B. (2009). Social image and the 50-50 norm: A theoretical and experimental analysis of audience effects. *Econometrica* 77, no. 5: 1607-1636.

[7] Andreoni, J., Callen, M., Jaffar, K.H., Khan, M.Y., & Sprenger, C.D. (2016). Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan. Working paper.

[8] Andreoni, J., & 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.

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

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

[11] Augenblick, N., Niederle, M., & Sprenger, C. (2015). Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics* 130 (3), 1067–1115.

[12] Bénabou, R., & Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review* 96, no. 5: 1652-1678.

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

[14] Bryan, C.J., Adams, G. S. & Monin, B. (2013). When cheating would make you a cheater: Implicating the self prevents unethical behavior. *Journal of Experimental Psychology: General* 142 (4), 1001–1005.

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

[16] Chowdhury, S. M., & Jeon, J-Y. (2014). Impure altruism or inequality aversion? An experimental investigation based on income effects. *Journal of Public Economics* 118, 143–150.

[17] Crumpler, H., & Grossman, P.J. (2008). An experimental test of warm glow giving. *Journal of Public Economics*, 92 (5), 1011–1021.

[18] Dana, J., Cain, D. M., & 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.

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

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

[21] Dreber, A., Fudenberg, D., Levine, D. K., & Rand, D. G. (2016). Self-Control, Social Preferences and the Effect of Delayed Payments. Working paper.

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

[23] Falk. A., & Zimmerman, F. (2016). Information Processing and Commitment. Working paper.

[24] Frederick, S, Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 351–401.

[25] Fudenberg, D., & Levine, D. (2006). A dual-self model of self-control. *American Economic Review* 96 (5), 1449–1476.

[26] Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review* 95 (1), 384–394.

[27] Gul, F. & Pesendorfer, W. (2001). Temptation and self-Control. *Econometrica* 69 (6), 1403–35.

[28] Halevy, Y. (2015). Time consistency: Stationarity and Time Invariance. *Econometrica* 83 (1), 335–352.

[29] Hanfling, O. (2008). How we trust one another. *Philosophy*, 83(02), 161–177.

[30] Harbaugh, W. T. (1998). The prestige motive for making charitable transfers. *American Economic Review* 88 (2), 277–282.

[31] Haushofer, J. & Shapiro, J. (2016). The short-term impact of unconditional cash transfers to the poor: Experimental Evidence from Kenya. *Quarterly Journal of Economics*, forthcoming.

[32] Kessler, J., Kivimaki, H. & Niederle, M. (2016). Generosity Over Time. Unpublished manuscript.

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

[34] Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 443-477.

[35] List, J.A. (2011). The market for charitable giving. *Journal of Economic Perspectives* 25 (2), 157–180.

[36] List, J.A., A.M. Shaikh, & Y. Xu (2016). Multiple Hypothesis Testing in Experimental Economics. NBER Working Paper 21875.

[37] Meyvis, T., A. Bennett & D.M. Oppenheimer (2010). Precommitment to Charity. *The Science of Giving: Experimental Approaches to the Study of Charity*. D. M. Oppenheimer & C.Y. Olivola (Eds.). New York: Taylor and Francis.

[38] O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *American Economic Review* 89 (1), 103–124.

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

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

[41] Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., & Greene, J. D. (2014b). Social heuristics shape intuitive cooperation. *Nature communications* 5, 3677.

[42] Recalde, M. P., Riedl, A., & Vesterlund, L. (2015). Error prone inference from response time: The case of intuitive generosity in public-good games. University of Pittsburgh Department of Economics WP, 15-004.

[43] Ribar, D. C., & Wilhelm, M. O. (2002). Altruistic and joy–of–giving motivations in charitable behavior. *Journal of Political Economy*, 110 (2), 425–457.

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

[45] Small, D. A., & Loewenstein, G. (2003). Helping a victim or helping the victim: Altruism and identifiability. *Journal of Risk and Uncertainty*, 26 (1), 5-16.

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

[47] Walton, G.M. & Banaji, M. R. (2004). Being what you say: The effect of essentialist linguistic labels on preferences. *Social Cognition* 22 (2), 193–213.

[48] Zellermayer, O. (1996). The pain of paying (Doctoral dissertation, Carnegie Mellon University).

Appendix
For Online Publication Only

Appendix A: Instructions

Experiment 1

Note: the text in square brackets was not read aloud. The instructions below are those that were used in the first set of sessions, with a show-up fee of $6 in the first week, and $20 in the second week. The instructions for the second set of sessions, that self-replicated our results with a show-up fee of $15 for each week, were the same except for the show-up fees.

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:]

- [Treatment 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.

- [Treatment 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.

- [Treatment 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.

- [Treatment 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.

Experiment 2

Your Donation Decision

In this study we will ask you to make two donation decisions, but only one of these two will end up being the decision that counts. One donation decision will be made today. Call this your week-1 donation decision. Your second donation decision will be made next week, when you return to the lab to complete this study. Call this your week-2 donation decision.

Here is how it works.

Week-1 donation decision

Today we will ask you whether you would like to donate $5 of your show up fee <u>for next week's session</u> 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," <u>$5 of your show up fee next week</u> will be donated. If you say NO, no donation will be made.

Week-2 donation decision

Next week, when you return to the lab to complete this study, you will have the opportunity to renew or revise your donation decision. In particular, next week you will be asked again whether you would like to donate $5 of your show up fee <u>for next week's session</u> to Give Directly. If you answer "YES, I'd like to donate $5 today," <u>$5 of your show up fee next week</u> will be donated. If you say NO, no donation will be made.

IMPORTANT: Only one of your decisions, either your week-1 or your week-2 donation decision, will be implemented. That is, only one decision will be the decision-that-counts. We will not use both! The most you will ever donate in this study is $5. The least you can donate is $0.

How will we decide whether your week-1 donation decision or your week-2 donation decision is the decision-that-counts?

Next week, <u>after</u> you make your week-2 donation decision, we will ask someone in the room to roll a 10-sided die to determine which decision is the decision-that-counts. All 10 numbers

on the die are equally likely. Based on your decision, there will be a 1 in 10 chance or a 9 in 10 chance that the decision-that-counts is your week-1 decision.

Today you will have two options to choose from:

A. Your **week-1 donation decision** will count with a **1 in 10 chance**, and so your week-2 donation decision will count with a 9 in 10 chance.

B. Your **week-1 donation decision** will count with a **9 in 10 chance**, and so your week-2 donation decision will count with a 1 in 10 chance.

If you chose Option A today, the following will occur. A volunteer will roll a 10-sided die and:

- Your **week-1** donation decision will be the decision-that-counts if number "1" is the outcome of the die roll.

- Your **week-2** donation decision will be the decision-that-counts if numbers "2", "3", "4", "5", "6", "7", "8", "9" or "10" are the outcome of the die roll.

If you choose Option B today, the following will occur. A volunteer will roll a 10-sided die and:

- Your **week-1** donation decision will be the decision-that-counts if numbers "2", "3", "4", "5", "6", "7", "8", "9" or "10" are the outcome of the die roll.

- Your **week-2** donation decision will be the decision-that-counts if number "1" is the outcome of the die roll.

In summary:

- Today you make a decision about donating $5 out of your show-up fee for next week's session to Give Directly. This decision will be carried out next week with a 1 in 10 or a 9 in 10 chance.

- Next week you will be asked again to make a decision about donating $5 out of your show up fee for next week's session to Give Directly. This decision will be carried out next week with a 9 in 10 or a 1 in 10 chance.

35

- Only one of these two decisions will be carried out.

- You make both donation decisions before you know which decision will be carried out.

- You decide today whether you would like your week-1 donation decision to count with a 1 in 10 chance (and so your week-2 donation decision will count with a 9 in 10 chance) or whether you would like your week 1 donation decision to count with a 9 in 10 chance (and so your week-2 donation decision will count with a 1 in 10 chance).

- After you have made your week-2 donation decision, a die will be rolled to determine whether your week-1 or your week-2 donation decision is the decision that counts.

Next you will be asked about your donation decision on the screens.

Remember: Your donation decision today could be the decision-that-counts so treat this decision as if it were the decision that will count.

Appendix B: Additional Analyses

B.1. Analysis of show-up rates

Table B.1 examines the determinants of the decision to show-up in week 2. We do not find that the treatment, or the decision to give or pledge within each treatment, or any individual characteristic is related to show-up in week 2.

Table B.1. Determinants of show-up in week 2

	(1)	(2)
	Show-up in Week 2	
Give-Later	0.027	0.023
	(0.030)	(0.030)
Pledge	-0.003	-0.005
	(0.036)	(0.038)
Pledge-or-Give-Now	0.039	0.024
	(0.038)	(0.040)
Give-Later * Gift in Week 1	-0.002	0.004
	(0.028)	(0.028)
Pledge * Pledge in Week 1	-0.010	-0.025
	(0.040)	(0.037)
Pledge-or-Give-Now * Pledge in Week 1	-0.047	-0.037
	(0.047)	(0.051)
Pledge-or-Give-Now * Gift in Week 1	-0.050	-0.047
	(0.035)	(0.032)
Female		-0.016
		(0.019)
Asian origin		-0.008
		(0.012)
Economics major		-0.033
		(0.021)
Senior student		-0.032
		(0.024)
Native English speaker		-0.018
		(0.023)
Nr. of correct questions in Cognitive Reflection Test (Frederick, 2005)		0.017
		(0.022)
Observations	692	689

Note: This table presents the average marginal effects (calculated at the means of all variables) from probit regressions on show-up in week 1. Column (1) presents the marginal effect from probit regressions on the treatment and its interaction with the decision made in week 1. Column (2) adds individual characteristics: female, which takes value one if the subject is a woman, zero otherwise; Asian origin, which takes value one if the subject's ethnicity is Asian, zero otherwise; Economics major, which takes value one if the subject is majoring in economics, zero otherwise; Senior student, which takes value one if the subject is in his or her senior year of studies or higher, zero otherwise; Native English speaker, which takes value one if the subject speaks English as a native language, zero otherwise; and Nr. of correct questions in the Cognitive Reflection Test (Frederick, 2005). A fixed effect for the show-up fee structure is included in all specifications. Robust standard errors, clustered at the session level, were used in each individual regression. ***,**,* indicates significance at the 1%, 5% and 10% levels, respectively.

B.2. Behavior of no-shows

Table B.2 displays week 1 decisions in Experiment 1, by show-up to week 2 of the experiment. Donation behavior in week 1 is not related to the decision to show-up in week 2.

Table B.2. Week 1 decision, by week 2 show-up rate

Treatment	Week 1 decision	Week 2 Show-up Yes	No	χ^2 test p-value
Give-Now	Give	31%	33%	0.846
Give-Later	Give	45%	45%	0.981
Pledge	Pledge	65%	63%	0.865
Pledge-or-Give-Now	Pledge	48%	39%	0.474
	Give	21%	17%	

B.3. Robustness tests: Results including no-shows

We examine the robustness of our results for Experiment 1 to the inclusion of participants who did not show-up for the second week of the experiment (no-shows). Table B.2. above already shows that behavior in week 1 of the experiment did not differ by participation in week 2. In Table B.3 we include participants who did not show up in week 2, in all treatments. More specifically, donation decisions were final in week 1 in the treatments Give-Now and Give-Later. Hence, these decisions are included in the data when no-shows are included for these treatments. In the treatments Pledge and Pledge-or-Give-Now individuals were expected to confirm or renege on their pledge in week 2. If they did not show up in week 2 we classify them as reneging on their pledge. This decision is guided by standard fundraising practices with pledging, where fundraisers contact individuals who pledged to confirm their pledge and receive their funds. Should they fail to contact them, they would not receive any donation from them. Table B.3. below displays donation and reneging behavior for each treatment. Results remain the same as those excluding no-shows, shown in the body of the paper.

Table B.3. Reneging and final donations, including no-shows

Treatment	Reneging	Donation rate
Give-Now	-	0.311
Give-Later	-	0.453
Pledge		
Pledge+Without thank-you	0.541	0.300
Pledge+Thank-you: Weak	0.425	0.375
Pledge-or-Give-Now		
Pledge-or-Give-Now+Without thank-you	0.731	0.337
Pledge-or-Give-Now+Thank-you: Weak	0.519	0.438
Pledge-or-Give-Now+Thank-you: Strong	0.510	0.442

B.4. Detailed Data on Commitment Demand

Table B.4 provides the detailed data on commitment demand and week 1 and week 2 decisions to donate in Experiment 2. These data are presented in Figure 3 in the body of the paper.

Table B.4. Giving and Commitment Demand.
Support for Figure 3

Week 1, Week 2 choices	Commitment choice			Subjects	
	No Commitment	Indifferent	Commitment	Number	Percent of all
Donate, Not Donate	64.9%	16.2%	18.9%	37	22.7%
Not Donate, Donate	34.8%	13.0%	52.2%	23	14.1%
Donate, Donate	29.3%	36.6%	34.1%	41	25.2%
Not Donate, Not Donate	19.4%	32.3%	48.4%	62	38.0%