

Altruism begets altruism: Nudging our way to a more virtuous society?*

Stephanie A. Heger[†]
School of Economics, The University of Sydney

Robert Slonim
School of Economics, The University of Sydney

August 8, 2019

ABSTRACT: Economic research examining social preferences over the past several decades has increasingly focused on better understanding and teasing apart distinct motives. While the research has primarily focused on short term behavior, this paper goes beyond this literature to better understand how a short-term intervention, a nudge, can affect subsequent behavior. Using a popular policy nudge, the default option, we show that its effect on the *choice* to be more altruistic “today” causes an increase in altruism “tomorrow”. We rule out that the nudge has a direct inter-temporal effect and instead show that our findings are consistent with a model of habit formation and moral consistency; that is, altruism begets altruism. Our local average treatment effect indicates that the nudge-induced giving in Round 1 of the experiment causes a 40 percentage point (or 200%) increase in the propensity to give in Round 2. Our findings suggest a way forward in promoting a more virtuous society.

*This project benefitted from many helpful comments and conversations, in particular, we thank Christine Exley, Lata Gangadharan, Tom Wilkening, Nina Xue and seminar participants at the University of Sydney and the University of Melbourne. Slonim gratefully acknowledges funding from ARC DP 150101307.

[†][Corresponding author] Email: saheger@gmail.com.

1 Introduction

Policy shapes society by encouraging socially desirable behavior. For example, in the United States, the government successfully incentivizes charitable giving by allowing individuals to deduct donations from their pre-taxed income. This policy can also have additional and often unaccounted for consequences: decreases in the after-tax price of giving increase charitable giving and are also associated with increases in other socially desirable behaviors, such as volunteerism (Feldman, 2010) and health (Yörük, 2014).¹²

A large body of research has also examined the factors that increase charitable giving, including the effects of price (Karlan and List, 2007), efficiency concerns Gneezy, Keenan, and Gneezy (2014); Exley (2015b), social pressure (List and Lucking-Reiley, 2002; Frey and Meier, 2004; Shang and Croson, 2009), and identity (Kessler and Milkman, 2016). However, there is far less research on the inter-temporal effects of these strategies. One exception is the literature on inter-temporal substitution in charitable giving, which examines how lowering the price to give (increasing the matching rate) today may also affect future giving (Cairns and Slonim, 2011; Castillo, Petrie, and Samek, 2017).

In this paper, we are interested in whether nudging an individual to make a charitable donation today will not only be effective today but, more importantly, also whether and how it affects altruism tomorrow.³ In other words, does altruism beget altruism?

Conceptually, the nudge we will examine, setting a default option, works “today” by decreasing the marginal psychological costs of choosing the desired behavior. Setting a default to nudge behavior has been found to successfully change the “today” decision in several contexts, such as how much to save for retirement (Benartzi and Thaler, 2007; Choi et al., 2003; Cronqvist and Thaler, 2004; Madrian and Shea, 2001) and joining an organ donor list (Kessler and Roth, 2012, 2014). A nudge may also have consequences on future behavior through two channels. First, the nudge “today” may have a direct inter-temporal effect on behavior “tomorrow”, which may work through salience to also make future pro-social behavior less psychologically costly. For example, Gneezy et al. (2012) report results from an experiment in which subjects who are randomly assigned to make a costly donation are more likely to behave honestly in a subsequent period than subjects who are randomly assigned to make a costless donation. Importantly, particularly in relation to our study, subjects in both the costly and costless treatment were forced to donate rather than having to choose whether to behave altruistically. Thus, while a direct impact through

¹Cappelen et al. (2017) also, find evidence of “unaccounted” for effects. In a large field experiment, they find that incentivizing subjects to go to the gym increases the likelihood of exercise, which in turn, increases the subjects’ academic performance.

²Policies can also be used to discourage undesirable behavior. For example, in the United States, the government discourages teenagers from drinking alcohol by placing a minimum age on alcohol consumption. Although this has decreased teen-drinking, it is also associated with an increase in teens’ use of marijuana (DiNardo and Lemieux, 2001).

³See Thaler and Sunstein (2003) and Sunstein and Thaler (2008) for a review of nudges.

salience is possible, Gneezy et al. (2012) prevents an indirect channel predicted by self-perception theory that we will explore here.

Specifically, a second (indirect) channel is that the nudge-influenced behavior today affects behavior in the future. In this paper, we hypothesize and then show that the nudge-induced increase in charitable giving today causes an increase in future charitable giving; that is, altruism begets altruism.

We hypothesize that identity drives the causal relationship between charitable giving. While there is an increased focus on the role of identity in economics (Akerlof and Kranton, 2005), including its role in charitable giving (Benjamin, Choi, and Fisher, 2010; Kessler and Milkman, 2016), Bem (1972)’s self-perception theory provides a framework for considering how identity might link altruism inter-temporally. Self-perception theory posits that individuals use past behavior and choices to make inferences about their own identity, which then inform future choices. Bénabou and Tirole (2011) formalize self-perception theory in economics and model individuals with imperfect memories of their identity but use their past actions to make inferences about their identity. This inference about their identity then provides a guide for current choices. Thus, self-perception theory predicts a path-dependency between moral actions over time.

Motivated by the history-dependence in actions modelled by Bénabou and Tirole (2011), we model our decision-maker’s utility at time t as dependent on his current choice of charitable giving and his $t - 1$ choice of altruism using a habit formation model (Pollak, 1970). Further, his $t - 1$ choice depends on his previous choices of altruism as well as on whether he is nudged towards altruism or selfishness at $t - 1$. Whether past charitable donations have a positive or negative effect on current altruistic behavior depends on how they enter into the individual’s utility function and is the parameter of interest in this paper.

Further, we incorporate identity into our model through past choices of altruism—the more the charitable donations an individual has given in the past, the more likely he is to strongly identify as altruistic. Our model also predicts that if altruism begets altruism, then this will be driven by individuals for whom altruism is a weak facet of their identity. This is consistent with Benabou and Tirole’s (2011) model, where individuals for whom altruism is a weakly held facet of their identity are predicted to behave in a more morally consistent manner. On the other hand, Benabou and Tirole’s (2011) model of self-perception theory also predicts that effective challenges to strongly-held aspects of identity “today” are met with contradictory responses “tomorrow”. Thus, depending on whether an individual has a weakly or strongly held conviction towards altruism, Bénabou and Tirole (2011) predicts either moral consistency (Nisan, 1985; Nisan and Horenczyk, 1990) or moral licensing (Khan and Dhar, 2006; Monin and Miller, 2001; Ploner and Regner, 2013; Sachdeva, Iliev, and Medin, 2009) (also see Blanken, van de Ven, and Zeelenberg (2015) and Mullen and Monin (2016) for a review of this literature).

To examine whether altruism begets altruism, we ran an online experiment. In our two main

treatments, we nudged individuals to either donate to charity or to keep the money for themselves by setting their default option to “donate” or to “keep”, respectively. Critically, subjects are either randomly nudged towards altruism or randomly nudged towards selfishness. To avoid donating, subjects in the Default Charity condition must opt-out of giving to charity; by contrast, subjects in the Default Cash condition must opt-in to giving to charity and opt-out of keeping cash (Round 1). Consistent with past research on default option nudges we find that our nudge positively impacts charitable giving behavior (Benartzi and Thaler, 2007; Choi et al., 2003; Cronqvist and Thaler, 2004; Madrian and Shea, 2001; Kessler and Roth, 2012, 2014). Specifically, we find that subjects in the Default Charity condition are twice as likely to donate in Round 1 than subjects in the Default Cash condition.

The critical part of the design is that at a later point in the experiment we ask subjects to make another donation to test whether initial altruistic behavior increases altruism in the future (Round 2). Directly motivated by our model and experimental design, we estimate a local average treatment effect and find that the nudge-induced increase in giving in Round 1 causes giving in Round 2 to increase by 200% or 40 percentage points. We also find that the nudge itself has no direct inter-temporal effect on giving in Round 2. Overall, our experiment shows that the nudge-induced altruism in Round 1 begets more altruism in Round 2,⁴ thus generating a virtuous cycle of altruism.

In addition and consistent with both our model and Benabou and Tirole’s (2011) model, individuals for whom altruism is a weakly held facet of their identity behave in a significantly more morally consistent manner. For these individuals, behaving altruistically in Round 1 causes an 83 percentage point (or 492%) increase in altruism in Round 2. Interestingly, we find that the Default Charity treatment does not differentially affect donation rates in Round 1 between subjects for whom altruism is a weak value and those for whom it is a strong facet of identity. Thus, we cannot attribute the differences in moral consistency between weak and strong identities to differences that stem from behavior in the first stage. However, self-perception theory offers a possible explanation; self-perception theory suggests that the altruistic behavior induced by the nudge is more informative for weak altruists than for strong altruists. Strong altruists have a richer history of donation behavior to draw from when making inferences about their identity to inform their Round 2 decision. On the other hand, weak altruists have a much sparser history that will make the Round 1 decision salient and easily recalled.

Our main result showing that the nudge-induced altruism in Round 1 causes altruism in Round 2 is based on estimates from a local average treatment effect. To support the validity of our econometric methods, we ran additional treatments in which we randomly assigned subjects to a default position in Round 1, but do not give them the choice to opt-out of their initial position;

⁴In fact, this finding is similar to the exclusion restriction assumption needed to estimate a local average treatment effect using instrumental variables (Angrist, Imbens, and Rubin, 1996)

that is, they are forced to make a donation or to keep the money in Round 1. Importantly, we find that these two groups donate at equal rates in Round 2, ruling out the possibility that the nudge has a direct inter-temporal effect, and thus providing further support that it is the choice to act altruistically, induced by our nudge, that causes the increase in altruism in Round 2. These additional treatments provide some evidence that the exclusion assumption for instrumental variables holds, providing additional support for the validity of our IV estimate.

Our contribution is thus twofold. First, we provide very strong evidence in favor of moral consistency; that is, we show that exogenously induced altruistic choices at $t - 1$ causes an increase in altruism at t . To claim this causal relationship, we show that the exclusion restriction assumption holds and that the nudge itself is not responsible for the inter-temporal effect, but rather the choice to act altruistically that the nudge induces at $t - 1$ causes the increase at t . Therein lies our second contribution—using experimental treatments to directly test that the theoretical assumptions behind our empirical test hold.

2 Experimental Design, Data & Hypotheses

In this section, we describe our experimental design and the data generated by the experiment. We also present a model of consumption choice, which motivates two competing hypotheses which we test in section 3.

2.1 Calibrating Preferences

We ran a pre-experimental calibration exercise to gauge the amount that must be donated to the chosen charity for the average subject to be indifferent to giving up \$1. The calibration exercise is important to set the default options such that some subjects will prefer to donate, while other subjects will prefer to keep cash for themselves. By finding a the median point of indifference between donating to charity and keeping cash for self, we can be confident that the nudge towards charity or the nudge towards keeping cash will be on the appropriate margins.

To do the calibration, we used the same charity, CARE, that we will use in the Round 1 decision of the experiment. This exercise follows the calibration exercise in Exley (2015a) and presents subjects with a multiple price list. On each line, they are asked whether they prefer to keep a \$1 and give \$0 to the charity or keep \$0 and give \$x to the charity, where $x \in \{\$0, \$0.1, \dots, \$3\}$. While Exley (2015a) uses a within-subject calibration, our calibration is taken as the median point of indifference across subjects, which was \$1 to self was utility-equivalent to \$1.50 to charity. This is how we chose the values in Round 1: subjects in the Default Cash condition were endowed with \$1 to keep for themselves and subjects in the Default Charity condition were endowed with making a \$1.50 donation to the charity. Subjects in this calibration exercise were excluded from

participating in any of the experimental conditions that follow.

2.2 Main Treatments

The main experiment consists of two Rounds. In Round 1, subjects were randomly endowed with \$1 cash (Default Cash condition) or endowed with a \$1.50 donation to the charity CARE (Default Charity condition). Figures A1a and A1b display what the subjects saw if they were assigned to the Default Cash and Default Charity treatments, respectively. After providing their endowment, we took two additional steps to facilitate a sense of ownership among subjects of their default position. First, we asked subjects in the Default Charity condition to list three ways the charity CARE might spend this money and we asked subjects in the Default Cash condition to list three ways they might spend their cash endowment. Second, we asked subjects to complete a set of unrelated filler questions. These filler questions created a period over which the subject had ownership of their default position (Strahilevitz and Loewenstein, 1998). Having subjects write about their endowment is a common technique in the psychology literature to increase the sense of ownership (will find refs) and elongating the time of having ownership of one's endowment has been shown to increase the endowment effect (Strahilevitz and Loewenstein, 1998). Moreover, while completing the filler questions,⁵ we reminded subjects of their default position by showing an image of their endowment to further reinforce the ownership of the default option they were given.

After completing the filler tasks, we asked subjects whether they would like to swap their position. Subjects assigned to the Default Cash treatment were asked if they wanted to give back their \$1 to make a \$1.50 donation to CARE while subjects assigned to the default donation treatment were asked if they wanted to not make the \$1.50 donation to get \$1 in cash. Figures A1c and A1d display the decisions faced by the subjects from the Default Cash and Default Charity treatments, respectively.

Next, we presented subjects with a multiple price list in which they had to choose one of 11 options. For each item, they could choose to add $\$X = (0, 0.10, 0.20 \dots 1.00)$ to their bonus and donate $\$2 \times (1-X)$ to Save the Children (see Table A1). For example, in the first option, subjects could choose to add \$1 to their bonus and donate \$0 to Save the Children, while in the last option, subjects could choose to add \$0 to their bonus and donate \$2 to Save the Children. Subjects had to make one choice from the list. We chose a new charity for the Round 2 decision to avoid a potential charity-specific wealth effect; that is, if some subjects donated to CARE in Round 1 (and others did not), then the marginal utilities of donating to CARE in Round 2 could differ by treatment assignment.

After completing the two Rounds of decisions, we asked a brief series of demographic questions

⁵Please see the full experimental protocol here to see the filler tasks the subjects performed.

as well as questions about their past charitable giving behavior. We summarize and discuss these statistics below in Table 1.

2.2.1 Additional Treatments

In addition to the two main treatments, we included three additional treatments. In our first and second control treatments, we again assign individuals to either the Default Charity or the Default Cash condition in Round 1, subjects complete filler tasks and are asked to make the same Round 2 donation decision. The main difference between these two treatments and the original treatments is that subjects are not given the opportunity to switch their Round 1 default position. In these two No Choice-No Info treatments, we randomly assign subjects to the Default Cash (No Choice-No Info) or the Default Charity (No Choice-No Info) and we never mention the opportunity or presence of an alternative position. Thus, subjects in the Default Charity (No Choice-No Info) condition all make a donation in Round 1 and are unaware of any option to have cash and subjects in the Default Cash (No Choice-No Info) all keep their cash in Round 1 and are unaware of any option to donate.

In sum, all subjects in the Default Charity (No Choice-No Info) condition will donate, and thus, importantly, they have no *active choice* to do so. Thus, comparing the Default Charity condition, which permits a Round 1 choice, to the Default Charity (No Choice-No Info), which does not permit a Round 1 choice, tests the critical prediction of self-perception theory for our main hypotheses. Specifically, self-perception theory argues that an active choice is an important component linking behavior and identity, noting that rejected alternatives reinforce the inferences an individual can make from their choice about their identity (Bem, 1972; Zanna, 1972). For example, an individual learns less about his altruistic identity if he is forced to donate rather than having had an active choice to keep the money for himself. In other words, the psychological cost associated with the donation in Round 1 is greater when the subject has the opportunity to keep the money than when he is forced to make a donation. The purpose of the No Choice-No Info treatments is thus to examine whether it is the *choice* of altruism in Round 1 that begets altruism in Round 2, rather than the nudge itself. This is critical for the interpretation of the results and also for establishing the exclusion restriction necessary to estimate a local average treatment effect.

In our third control treatment (henceforth: Round 2 Only), subjects do not make a Round 1 decision and instead begin the experiment with the filler tasks and are then asked to make a Round 2 donation decision that is identical to the original treatments. The purpose of the Round 2 Only treatment is twofold. First, if we find a difference between the Default Charity and Default Cash, we can compare Round 2 behavior in these treatments to the behavior of subjects in the Round 2 Only condition to see which of those treatments are more similar to the Round 2 Only treatment. If one of the default treatments is closer to the Round 2 Only behavior, then this will tell that the other default treatment was the treatment that had the bigger impact on Round 2

choices. Second and related, we can examine whether subjects who have the opportunity to give twice (i.e., subjects in the Default Charity and Default Cash) make more total donations than subjects who are only asked to give once (i.e., Round 2 Only); that is, we can examine whether the second ask does not crowd-out donations.

2.3 Data

The data from our experiment come from Amazon’s Mechanical Turk (Mturk). The experiment was completed by 1020 Mturk workers from the United States. Table 1 displays the summary statistics for the subjects in our experiment. Approximately half of the subjects are female, the majority work full-time and 77% have donated to charity at least once in the last year.

To determine subject’s conviction towards altruism, we asked them whether they had donated money to a charity 0, 1, 2, 3 or 4 or more times in the past year (not including the donation made during the experiment). On average, subjects donated 2.3 times in the past 12 months. Using this variable, we classify subjects as having a strong conviction towards altruism if they indicated that they have given 4 or more times in the past year.

TABLE 1: SUMMARY STATISTICS

	All	Treatment Conditions				Round 2 Only
		Default Charity	Default Cash	Default Charity (No Choice-No Info)	Default Cash	
No. of Donations in Past Year	2.35 (1.61)	2.49 (1.62)	2.15 (1.66)	2.45 (1.59)	2.22 (1.59)	2.47 (1.58)
Altruism Strongly Held Value Past donations ≥ 4	.34 (.47)	.38 (.49)	.31 (.46)	.33 (.47)	.28 (.45)	.38 (.49)
No Donation in Past Year	.23 (.42)	.21 (.40)	.29 (.45)	.21 (.41)	.24 (.43)	.21 (.41)
Have Donated in Past Year	.77 (.42)	.80 (.40)	.71 (.45)	.79 (.41)	.76 (.43)	.79 (.41)
Female	.53 (.50)	.52 (.50)	.54 (.50)	.54 (.50)	.54 (.50)	.54 (.50)
Age	38.35 (11.80)	40.51 (11.82)	38.38 (11.56)	37.54 (12.16)	37.03 (12.06)	38.36 (11.24)
Unemployed	.06 (.24)	.05 (.22)	.09 (.29)	.03 (.17)	.07 (.26)	.08 (.26)
Employed full-time	.62 (.49)	.64 (.48)	.58 (.49)	.63 (.48)	.64 (.48)	.59 (.49)
Employed part-time	.16 (.37)	.14 (.35)	.18 (.38)	.19 (.39)	.14 (.35)	.17 (.37)
Retired	.04 (.18)	.05 (.21)	.04 (.20)	.04 (.20)	.03 (.16)	.02 (.14)
Income < \$10,000	.05 (.22)	.07 (.25)	.05 (.23)	.04 (.20)	.04 (.19)	.05 (.22)
Income > \$150,000	.03 (.18)	.04 (.19)	.02 (.15)	.04 (.21)	.02 (.13)	.05 (.21)
Observations	1020	191	224	224	182	199

Means reported with standard deviations in parentheses.

2.4 Model, Hypotheses and Empirical Strategy

Next, we turn to modeling the choice to donate at t , given previous donation choices, and the main question of our paper: does altruism beget altruism? To formalize this question, we consider an individual who has preferences over two goods at time t , private consumption (c_t) and charitable giving (A_t). The individual's preferences can be represented by a utility function with the following form,

$$U(c, A) = u(c_t, c_{t-1}(\Theta_c)) + \alpha v(A_t, A_{t-1}(\Theta_A)) \quad (1)$$

where Θ_c and Θ_A represent a composite of private consumption and charitable giving up to and including time $t - 2$, respectively. Thus, today's utility depends on the choices the individual makes today as well as all past choices. The parameter $\alpha \in [0, 1]$ governs the intensity of the individual's preference for altruism and warm glow. The functions $u(\cdot)$ and $v(\cdot)$ are concave in consumption and donations to charity, respectively. A subject solves the following equation at time t

$$\max_{c_t, A_t} U(c_t, A_t \mid \bar{c}, \bar{A}) = \max_{c_t, A_t} u(c_t + \gamma_c c_{t-1}(\Theta_c)) + \alpha v(A_t + \gamma_A A_{t-1}(\Theta_A)) \text{ subject to } I = c_t + p \times A_t \quad (2)$$

where the parameter γ_c and $\gamma_A \in \mathbb{R}$ represent the intensity of the past consumption ($c_{t-1}(\Theta_c)$, $A_{t-1}(\Theta_A)$) on today's utility and will pin down whether there are negative, positive or no spillovers. I is income and p is the relative price of making a donation. We want to compare the optimal choices at time t of individuals nudged towards altruism versus subjects nudged towards selfishness at time $t - 1$. Let $A_t(Z)$ and $A_{t-1}(\Theta_A, Z)$ represent the choices at t and $t - 1$, respectively, for an individual who receives nudge $Z \in 0, 1$, where $Z = 1$ indicates the subject was nudged towards altruism and $Z = 0$ indicates the individual was nudged towards selfishness. From the first order conditions we find that

$$A_t(Z = 0) + \gamma_A A_{t-1}(\Theta_A, Z = 0) = A_t(Z = 1) + \gamma_A A_{t-1}(\Theta_A, Z = 1) \quad (3)$$

We assume that $\frac{\partial A_{t-1}}{\partial \Theta_A} \Big|_{Z=1} \geq \frac{\partial A_{t-1}}{\partial \Theta_A} \Big|_{Z=0}$. Rearranging and taking expectations of equation 3, we obtain

$$\frac{E[A_t \mid Z = 1] - E[A_t \mid Z = 0]}{E[A_{t-1} \mid \Theta_A, Z = 1] - E[A_{t-1} \mid \Theta_A, Z = 0]} = (-1)\gamma_A \quad (4)$$

The left-hand-side of equation 4 is the equation for an instrumental variable estimand, β^{IV} . Thus, we propose to test for positive or negative spillovers by estimating the local average treatment effect (Imbens and Angrist, 1994) using instrumental variables (Angrist, Imbens, and Rubin, 1996).

Our identification strategy relies on three assumptions. First, the instrument, Z , is randomly assigned. We satisfy this assumption in our experimental design. Second, the effect of the instrument, Z , must be monotonic in that a subject in the Default Charity condition must be at least as likely to donate in Round 1 than he would have been had he been assigned to the Default Cash condition. Third, the exclusion restriction states that the instrument only affects outcome A_t through A_{t-1} (i.e., Round 1 donation behavior) and does not directly affect outcomes, A_t . Given our research question and experimental design, this assumption is the hardest to justify without some evidence. However, in our control treatments described in Section 2.2.1, we removed the *choice* to donate in Round 1. If there are no treatment differences in Round 2 behavior when there is no active choice in Round 1 (i.e., the No Choice treatments), then we take this as evidence that any differences in Round 2 when there is an active choice (i.e., the original treatments) can be attributed to the treatment-induced change in behavior in Round 1 rather than solely to the treatment itself. This suggests that it is the choice of donating in Round 1, which is influenced by the subject’s default position, rather than the nudge (i.e., the instrument, Z) that affects Round 2 donation decisions (i.e., outcomes, A_t).

Our first hypothesis concerns the denominator of equation 4 and is a test of the first stage of our IV estimate. In particular, we hypothesize that, on average, subjects in the Default Charity condition will be more likely to donate in Round 1 than subjects in the Default Cash condition; that is, $E[A_{t-1} | \Theta_A, Z = 1] - E[A_{t-1} | \Theta_A, Z = 0] > 0$.

Hypothesis 1. *Default Option Hypothesis:* *Participants who are endowed with a donation (i.e., defaulted into donating) are more likely to donate in Round 1 than participants who are endowed with cash (i.e., defaulted into keeping cash).*

Our second hypothesis uses data from the additional treatments, in particular, the No Choice-No Info treatments, to test the validity of the Exclusion Restriction assumption needed for estimating and interpreting β^{IV} .

Hypothesis 2. *Exclusion Restriction Hypothesis:* *The default option treatment Z does not directly affect the decision to donate in Round 2. Instead, any effect of Z on Round 2 donation choices operates solely through the choice to donate in Round 1. Subjects in the Default Charity (No Choice-No Info) will be equally as likely as subjects in the Default Cash (No Choice-No Info) to donate in Round 2.*

Finally, we turn to the main hypotheses about the direction of the behavioral spillovers. Positive behavioral spillovers imply that $E[A_t | Z = 1] - E[A_t | Z = 0] > 0$ and therefore, given Hypothesis 1, $\gamma_A < 0$. We interpret a positive behavioral spill-over as moral consistency since $\gamma_A < 0$ implies that the Default Charity condition exogenously increases altruism in Round 1 and that this nudge-induced increase in altruism in Round 1 causes an increase in altruism in Round 2.

Hypothesis 3. Moral Consistency Hypothesis: $\beta^{IV} > 0$, implying that $\gamma_A < 0$ which means that an increased propensity to donate in Round 1 will:

- (i) increase the propensity to donate in Round 2;
- (ii) increase the amount donated in Round 2.

Conversely, if $\beta^{IV} < 0$ then $\gamma_A > 0$ which means the nudge-induced altruism in Round 1 causes a decrease in altruism in Round 2. We interpret this as evidence consistent with moral licensing.⁶

To econometrically analyze the experimental data, we estimate a two-stage least squares regression, where we first estimate the effect of the treatment assignment, Z_i , on Round 1 donation behavior, $A_{i,t-1}$. We then use the predicted values of Round 1 donation behavior, $\widehat{A_{i,t-1}}$, to estimate the second stage to obtain the causal effect of donating in Round 1 on donating in Round 2, Y_i . The interpretation of the coefficient, β^{IV} , is the change in Round 2 donation rates that are caused by the *treatment-induced* donation behavior in Round 1.

$$A_{i,t} = \beta_0 + \beta^{IV} \widehat{A_{i,t-1}} + \varepsilon_i, \quad (5)$$

We then examine how identity affects the causal relationship between altruism at t and $t - 1$. We draw from self-perception theory and posit that identity is inferred from past choices. Thus, the composite of past charitable giving, Θ_A , serves as a proxy for the facet of identity related to altruism.

Weak Identity Towards Altruism It is straightforward to show that if $\gamma_A < 0$, then γ_A is *increasing* (towards 0) in the strength of the individual’s altruistic identity (Θ_A). On the other hand, if $\gamma_A > 0$, then γ_A is *decreasing* (towards 0) in the strength of the individual’s altruistic identity (Θ_A). Thus, Hypothesis 4 states that the magnitude of the local average treatment effect will be greater for those with a weak identity than for those with a strong identity towards altruism.

Benabou & Tirole’s (2011) model, also drawing heavily from self-perception theory, predicts that when weakly-held values are encouraged, individuals respond in a confirmatory way (i.e., morally consistent), as the value becomes more salient to the individual. In other words, as stated in Hypothesis 5, individuals for whom altruism is a weak facet of their identity will behave in a morally consistent way in the future, when nudged towards altruism today. Thus, the predictions from our model and the Bénabou and Tirole (2011) model about individuals with weak identities towards altruism are similar when $\gamma_A < 0$.

⁶**Moral Licensing and Negative Spill-over Hypothesis:** $\beta^{IV} < 0$, implying that $\gamma_A > 0$ which means that an increased propensity to donate in Round 1 will:

- (i) decrease the propensity to donate in Round 2;
- (ii) decrease the amount donated in Round 2.

Strong Identity Towards Altruism By contrast, Bénabou and Tirole (2011) predict that when strongly-held convictions are challenged, individuals will be more likely to respond in a contradictory way to the challenge to restore their self-image. This means that for those individuals who have a strong identity towards altruism but are nudged towards selfishness (i.e., the Default Cash condition), Bénabou and Tirole (2011) predicts that individuals will respond by being more altruistic in the future. We formally state this in Hypothesis 6.

To test these hypotheses, we will use a similar specification from equation 5. One change is that we will interact our endogenous regressor (A_{t-1}) with the strength of conviction towards altruism, either weak ($\Theta_A = 0$) or strong ($\Theta_A = 1$), and instrument for Round 1 donation behavior using the assignment to the Default Charity treatment interacted with the strength of the conviction. Our specification for this hypothesis is therefore given by

$$A_{i,2} = \delta_0 + \delta_1^{IV} \widehat{A}_{i,1} \times \mathbf{1}[\Theta_A = 0] + \delta_2^{IV} \widehat{A}_{i,1} \times \mathbf{1}[\Theta_A = 1] + \varepsilon_i, \quad (6)$$

Similarly, to estimate the causal effect of keeping the money in Round 1 ($c_{i,1}$) on the likelihood of also keeping the money in Round 2 ($c_{i,2}$) we instrument for keeping the money in Round 1 using the assignment to the Default Cash condition. We use a similar interaction as in equation 6.

$$c_{i,2} = \lambda_0 + \lambda_1^{IV} \widehat{c}_{i,1} \times \mathbf{1}[\Theta_A = 0] + \lambda_2^{IV} \widehat{c}_{i,1} \times \mathbf{1}[\Theta_A = 1] + \varepsilon_i, \quad (7)$$

Hypothesis 4. *Altruism as a Weak Facet of Identity I:* if $\gamma_A < 0$ ($\gamma_A > 0$), then individuals who hold altruism as a weak facet of their identity will behave more morally consistent (balanced) than individuals who hold altruism as a strong facet of their identity; that is, $|\delta_1^{IV}| > |\delta_2^{IV}|$.

Hypothesis 5. *Altruism as a Weak Facet of Identity II:*

$\delta_1^{IV} > 0$: individuals who hold altruism as a weak facet of their identity will behave morally consistently.

Hypothesis 6. *Altruism as a Strong Facet of Identity:*

(i) $\lambda_2^{IV} < 0$: individuals who hold altruism as a strongly-held facet of their identity will respond in a contradictory way to a nudge towards selfishness.

3 Results

In Section 3.1, we present the main results as outlined by Hypotheses 1-6. In Section 3.2, we present additional findings.

3.1 Main Results

In this section, we outline our main results based on our hypotheses from Section 2.4. Taken together, our set of results is consistent with moral consistency; that is, exogenously-induced altruism in one period causes an increase in altruism in a subsequent period. In other words, altruism begets altruism. To establish this result, we must first demonstrate that we have a strong instrument (result 1) and, importantly, that the exclusion restriction holds (result 2). We then establish our main result, that altruism begets altruism, in result 3. For each result, we begin with a statement of the result, followed by the evidence to support it.

3.1.1 Round 1 Decisions

Result 1. *We find evidence of an endowment effect for donation behavior. Subjects who must opt-out of the donation are significantly more likely to donate in Round 1 than subjects who must opt-in.*


Figure 1a show that subjects assigned to the Default Charity treatment are 19 percentage points more likely (more than a 100 percent increase) to donate to charity in Round 1 than subjects assigned to the Default Cash treatment (t-test: p-value<.0001). In Figure 1b we look at the effect of the default option nudge by subjects' number of charitable donations in the past 12 months. In general, we find that the Default Charity condition significantly increases donation rates regardless of the subjects' past donation history. For example, subjects with strongly-held (weakly-held) views towards altruism are 23 (15) percentage points more likely to give under the Default Charity than Default Cash condition (t-test: p-value=.003 and p-value=.002, respectively). Moreover, in a regression framework, we show that the Default Charity nudge increases the propensity to give in Round 1 by 18 percentage points while having a subject with a strong conviction towards altruism, relative to a weak one, is associated with an 8 percentage point increase in giving in Round 1. The difference between the effect of the nudge and the effect of convictions is not significant (p-value=.15).

3.1.2 Round 2 Decisions

Result 2. *We find that the treatment itself has no direct inter-temporal effect; that is, subjects in the Default Charity (No Choice) do not behave more altruistically in Round 2 than subjects in the Default Cash (No Choice).*

Table 2 shows the relationship between donation behavior in Round 2 and treatment assignment. Columns (1), (3) and (5) of Table 2 show that subjects in the Default Charity condition are marginally more likely to donate and make larger donations in Round 2 than subjects in the Default Cash condition. Further, columns (2), (4) and (6) show that this difference is driven by

FIGURE 1: FIRST-STAGE: ROUND 1 DONATION RATES



Average donation rates in Round 1 by treatment assignment with 95% confidence intervals.

subjects for whom altruism is a weakly-held value. For example, subjects who hold altruism as a weakly-held value are 13 percentage points more likely to donate in Round 2 in the Default Charity condition than in the Default Cash condition. Columns (2), (4) and (6) show no differences between the Default Charity and the Default Cash conditions among subjects for whom altruism is a strongly-held value.

However, the treatment effects in Round 2 do not indicate whether the higher donation rates in Round 2 in the Default Charity condition are due to a direct inter-temporal effect of the treatment or whether it is the increase in giving caused by the treatment in Round 1 that causes an increase in Round 2 giving. To rule out the former, we examine the average donation rates and amounts for subjects in the No Choice-No Info treatments, where subjects do not make an active choice in Round 1.⁷ The first F -test in Table 2 shows that subjects in the Default Charity (No Choice-No Info) are not more likely to donate, nor donate more, than subjects in the Default Cash (No Choice-No Info). Further, Columns (2), (4) and (6) show that these results are robust regardless of whether subjects conviction to altruism is weak or strong. In sum, we interpret this as evidence that our treatments have no direct effect on Round 2 behavior and that the exclusion restriction hypothesis holds.

Gneezy et al. (2012) find that the (monetary) cost of a pro-social act is an important component in shaping self-image. We extend this finding to also include cognitive costs—in the Default Charity (No Choice-No Info), where there are no choice alternatives to donating and thus the

⁷In the No Choice-No Info treatments we vary two components with respect to the main treatments. First, we vary whether subjects' awareness of another alternative (No Info) and second, we remove the choice. To test whether the removal of the information or thinking about a choice could explain the differences between our main treatments and the No Choice-No Info conditions, we ran four additional treatments, two No Choice-Info treatments and two Hypothetical Choice treatments, thus rejecting the possibility that differences in information or the consideration of making a choice could explain the different results between our main treatments and the the No Choice-No Info conditions. The results looking at these four additional treatment conditions are available upon request.

decision to donate is costless, there are no significant spillovers into future altruism. On the other hand, in the main Default Charity condition, where there is a choice alternative and thus the decision to donate is more cognitively costly, there are significant spillovers.

TABLE 2: TESTS FOR EXCLUSION RESTRICTION: THE IMPORTANCE OF CHOICE IN ROUND 1

	Propensity to Donate		Donation Amount			
	(1)	(2)	(3)	(4)	(5)	(6)
Default Charity	0.08 (0.05)	0.13** (0.06)	0.11 (0.07)	0.15* (0.08)	-0.26* (0.14)	0.51** (0.23)
Default Charity (No Choice-No Info)	0.15*** (0.05)	0.21*** (0.05)	0.06 (0.06)	0.1 (0.07)	-0.24* (0.13)	0.57*** (0.22)
Default Cash (No Choice-No Info)	0.21*** (0.05)	0.24*** (0.06)	0.08 (0.06)	0.08 (0.07)	-0.13 (0.14)	0.57** (0.23)
Default Charity \times Strong Value	.	-0.17* (0.1)	.	-0.14 (0.15)	.	-0.53 (0.38)
Default Charity (No Choice-No Info) \times Strong Value	.	-0.19* (0.1)	.	-0.13 (0.14)	.	-0.62* (0.37)
Default Cash (No Choice-No Info) \times Strong Value	.	-0.09 (0.11)	.	0.02 (0.15)	.	-0.24 (0.39)
Strong Value	.	0.2*** (0.07)	.	0.24** (0.1)	.	0.78*** (0.27)
Observations	821	821	821	821	821	821
R^2	0.02	0.04	0.004	0.02	.	.
Omitted Group			Default Cash			
Mean probability, Default Cash	.32	.32	.37	.37	.37	.37
F-Tests						
[1] Default Charity (No Choice-No Info) = Default Cash (No Choice-No Info)	1.18	.19	.11	.06	.39	0
[2] Default Charity = Default Charity (No Choice-No Info) – Default Cash (No Choice-No Info)	3.61*	3.41*	2.07	1.54	.40	2.66

Columns (1) - (4), OLS regression estimates. Columns (5) & (6) Tobit regression estimates. Robust standard errors in parentheses and *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

Moral Consistency

Result 3. *Consistent with Hypothesis 3, we find evidence of moral consistency, i.e., $\beta^{IV} > 0$, implying that $\gamma_A < 0$; that is, increased giving in Round 1 increases giving in Round 2. In particular, giving in Round 1 increased the probability of giving in Round 2 by 200% (40 percentage points). Further, giving in Round 1 increases the amount given in Round 2 by \$.59.*

We estimate the effect that the nudge-induced increase in giving in Round 1 has on charitable giving in Round 2 using the instrumental variable approach discussed in Section 2.4. The local average treatment effect estimates are presented in Table 3. In Panel A, we estimate equations 5 and 6 and in Panel B we estimate equation 7.

Column (1) of Panel A indicates that giving in Round 1 causes a 40 percentage point (200% increase above the baseline) increase in the propensity to give in Round 2 (p-value=.071). In columns (3)-(6) the dependent variable is donation amount in Round 2. Column (3) indicates that giving in Round 1 causes subjects to increase their giving by \$0.59 (200%) in Round 2 (p-value=.068). In sum, altruism begets altruism.

Result 4. *Consistent with our model and hypothesis 4, we find that subjects for whom altruism is a weak facet of their identity are significantly more morally consistent than those individuals for whom altruism is a strong facet of their identity.*

Result 5. *Consistent with Bénabou and Tirole (2011) and hypothesis 5, we find that subjects for whom altruism is a weakly-held value behave in a morally consistent manner.*

In columns (2), (4) & (6), we investigate the differential response in Round 2 of subjects with strongly-held versus weakly-held values towards altruism. Panel A shows, consistent with Bénabou and Tirole (2011), that subjects for whom altruism is a weakly-held conviction respond in a confirmatory way in Round 2 to their nudge-induced behavior in Round 1; that is, they demonstrate *more* moral consistency than subjects for whom altruism is a strongly-held value. One potential concern is that an individual who has only given to charity once or twice in the past year, but gave a large sum of money, would be classified as having a weak identity towards altruism under our definition. To address this potential problem, in Table A2 in Appendix A, we show that subjects who have not given in the past year (i.e., weak identity) are significantly more morally consistent than subjects who have given 4 or more times in the past year (i.e., strong identity).

Immoral Consistency In Panel B, we estimate equation 7 to examine whether there is evidence of immoral consistency; that is, does keeping the cash in Round 1 cause an increase in keeping the cash in Round 2. We do not find evidence consistent with immoral consistency on the extensive margin (columns (1) & (2)), but columns (3)-(6) show that keeping the cash in Round 1 causes subjects to keep more cash in Round 2.

Result 6. *We find no evidence that subjects who have a strong identity towards altruism behave in a morally balanced manner when nudged towards selfishness.*

Columns (2), (4) & (6) test for the second part of the Bénabou and Tirole (2011) hypothesis, which states that subjects who are nudged away from a strongly-held value will respond in a contradictory manner. Thus, we hypothesized that subjects for whom altruism is a strongly-held value, but are nudged towards selfishness, would be altruistic in Round 2. However, we do not find support for this hypothesis.

The results in Table 3 suggest that nudging virtuous behavior “today” may promote virtuous behavior “tomorrow”, particularly among those individuals who have been less virtuous in the past. In other words, the nudge successfully crowds people into giving in Round 2, who would likely not have given in Round 2, by nudging them to give in Round 1.

TABLE 3: LOCAL AVERAGE TREATMENT EFFECTS: ROUND 2 DONATION RATES & AMOUNTS

<i>Panel A: Moral Consistency</i>						
	Propensity to Donate			Donation Amount		
	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{A}_{i,1}$	0.41* (0.23)	.	0.59* (0.32)	.	1.77 (1.25)	.
$\widehat{A}_{i,1} \times StrongValue$.	-0.18 (0.39)	.	0.04 (0.53)	.	-0.77 (1.79)
$\widehat{A}_{i,1} \times WeakValue$.	0.83** (0.36)	.	0.95** (0.49)	.	3.81* (2.06)
Strong Value	.	0.36** (0.15)	.	0.37* (0.21)	.	1.69** (0.77)
Constant	0.25*** (0.06)	0.14* (0.08)	0.27*** (0.08)	0.15 (0.1)	-1.20*** (0.36)	-1.80*** (0.51)
Observations	415	415	415	415	415	415
R^2	0.2	0.03	0.21	0.12	.	.
χ^2 test						
$\widehat{A}_{i,1} \times Strong = \widehat{A}_{i,1} \times Weak$	3.59*		1.58		2.80*	
<i>Panel B: Immoral Consistency</i>						
	Propensity to Keep			Keep Amount		
	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{c}_{i,1}$	0.15 (0.16)	.	0.3* (0.16)	.	0.32* (0.18)	.
$\widehat{c}_{i,1} \times StrongValue$.	0.17 (0.24)	.	0.02 (0.27)	.	-0.38 (0.9)
$\widehat{c}_{i,1} \times WeakValue$.	0.11 (0.23)	.	0.48** (0.24)	.	1.91* (1.03)
Strong Value	.	-0.09 (0.25)	.	0.27 (0.27)	.	1.45 (1.02)
Constant	0.78*** (0.12)	0.82*** (0.18)	0.57*** (0.12)	0.45** (0.19)	0.54*** (0.14)	-0.007 (0.8)
Observations	415	415	415	415	415	415
R^2	0.08	0.09	0.21	0.12	.	.
χ^2 test						
$\widehat{c}_{i,1} \times Strong = \widehat{c}_{i,1} \times Weak$.03		1.58		2.80*	

Columns (1)-(4), OLS regression estimates. Column (5) & (6) Tobit regression estimates. Robust standard errors in parentheses and *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

3.2 Additional Findings: Multiple Donation Asks & Giving Behavior

In this section, we show that our nudge towards altruism, and moral consistency, in particular, helps to overcome decreases in giving that are typically associated with ask fatigue and multiple donation solicitations. Because we find that altruism begets altruism, this implies that giving in Round 1 begets more giving in Round 2. However, how do our treatment subjects compare to those subjects who are only asked to donate once? We make this comparison in Table 4, where the omitted group is the Control condition. We show that subjects in the Default Charity and the Default Cash condition who are asked to give in Round 1 and Round 2, do not give significantly less in Round 2 than subjects in the Control condition, who are only asked to give in Round 2. Further, this equivalence in Round 2 giving means that total giving is greater for subjects in the Default Charity and Default Cash conditions than for subjects in the Control.

TABLE 4: TOTAL GIVING

	Donation Amount in Round 2		Total Donation Amount	
	(1)	(2)	(3)	(4)
Default Charity	0.07 (0.07)	0.12 (0.08)	0.59*** (0.1)	0.57*** (0.11)
Default Cash	-0.04 (0.06)	-0.03 (0.07)	0.21** (0.08)	0.2** (0.09)
Default Charity \times Strong Value	.	-0.14 (0.14)	.	0.06 (0.21)
Default Cash \times Strong Value	.	-0.0006 (0.14)	.	0.08 (0.18)
Strong Value	.	0.24*** (0.09)	.	0.24*** (0.09)
Constant	0.41*** (0.04)	0.32*** (0.05)	0.41*** (0.04)	0.32*** (0.05)
Observations	614	614	614	614
R^2	0.005	0.03	0.06	0.07

OLS regression estimates. Robust standard errors in parentheses and *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

4 Conclusion

If altruism begets altruism, then a nudge towards pro-sociality may provide previously unaccounted for benefits in various arenas, including governmental policy and corporate culture. For example, tax policies that provide subsidies for individuals who give to charity may increase the rates of charitable giving (Gruber, 2004; Yörük, 2013), but also have the added benefit of increasing individuals' altruistic identities and thus leading to additional altruism. In an age when corporate culture, particularly that of the banking culture is highly scrutinized for its corruption and immorality (Cohn, Fehr, and Maréchal, 2014), a simple nudging of employees towards cooperative

behavior may reorient the corporate culture towards inclusivity and pro-sociality. Of course, the lasting effects of such nudges remains an open question.

References

- Akerlof, George A and Rachel E Kranton. 2005. "Identity and the Economics of Organizations." *Journal of Economic perspectives* 19 (1):9–32.
- Angrist, Joshua D, Guido W Imbens, and Donald B Rubin. 1996. "Identification of causal effects using instrumental variables." *Journal of the American statistical Association* 91 (434):444–455.
- Bem, Daryl J. 1972. "Self-perception theory." In *Advances in experimental social psychology*, vol. 6. Elsevier, 1–62.
- Bénabou, Roland and Jean Tirole. 2011. "Identity, morals, and taboos: Beliefs as assets." *The Quarterly Journal of Economics* 126 (2):805–855.
- Benartzi, Shlomo and Richard H Thaler. 2007. "Heuristics and biases in retirement savings behavior." *The journal of economic perspectives* 21 (3):81–104.
- Benjamin, Daniel J, James J Choi, and Geoffrey W Fisher. 2010. "Religious identity and economic behavior." Tech. rep., National Bureau of Economic Research.
- Blanken, Irene, Niels van de Ven, and Marcel Zeelenberg. 2015. "A meta-analytic review of moral licensing." *Personality and Social Psychology Bulletin* 41 (4):540–558.
- Cairns, Jason and Robert Slonim. 2011. "Substitution effects across charitable donations." *Economics Letters* 111 (2):173–175.
- Cappelen, Alexander W, Gary Charness, Mathias Ekström, Uri Gneezy, and Bertil Tungodden. 2017. "Exercise improves academic performance." .
- Castillo, Marco, Ragan Petrie, and Anya Samek. 2017. "Time to Give: A Field Experiment on Intertemporal Charitable Giving." .
- Choi, James J, David Laibson, Brigitte C Madrian, and Andrew Metrick. 2003. "Optimal defaults." *The American Economic Review* 93 (2):180–185.
- Cohn, Alain, Ernst Fehr, and Michel André Maréchal. 2014. "Business culture and dishonesty in the banking industry." *Nature* 516 (7529):86.
- Cronqvist, Henrik and Richard H Thaler. 2004. "Design choices in privatized social-security systems: Learning from the Swedish experience." *The American Economic Review* 94 (2):424–428.
- DiNardo, John and Thomas Lemieux. 2001. "Alcohol, marijuana, and American youth: the unintended consequences of government regulation." *Journal of health economics* 20 (6):991–1010.
- Exley, Christine L. 2015a. "Excusing selfishness in charitable giving: The role of risk." *The Review of Economic Studies* 83 (2):587–628.
- . 2015b. "Using charity performance metrics as an excuse not to give." Tech. rep., Working paper.
- Feldman, Naomi E. 2010. "Time is money: Choosing between charitable activities." *American Economic Journal: Economic Policy* 2 (1):103–30.
- Frey, Bruno S and Stephan Meier. 2004. "Social comparisons and pro-social behavior: Testing" conditional cooperation" in a field experiment." *American Economic Review* 94 (5):1717–1722.

- Gneezy, Ayelet, Alex Imas, Amber Brown, Leif D Nelson, and Michael I Norton. 2012. "Paying to be nice: Consistency and costly prosocial behavior." *Management Science* 58 (1):179–187.
- Gneezy, Uri, Elizabeth A Keenan, and Ayelet Gneezy. 2014. "Avoiding overhead aversion in charity." *Science* 346 (6209):632–635.
- Gruber, Jonathan. 2004. "Pay or pray? The impact of charitable subsidies on religious attendance." *Journal of Public Economics* 88 (12):2635–2655.
- Imbens, Guido W and Joshua D Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62 (2):467–475.
- Karlan, Dean and John A List. 2007. "Does price matter in charitable giving? Evidence from a large-scale natural field experiment." *American Economic Review* 97 (5):1774–1793.
- Kessler, Judd B and Katherine L Milkman. 2016. "Identity in charitable giving." *Management Science* 64 (2):845–859.
- Kessler, Judd B and Alvin E Roth. 2012. "Organ allocation policy and the decision to donate." *The American Economic Review* 102 (5):2018–2047.
- . 2014. "Getting more organs for transplantation." *The American Economic Review* 104 (5):425–430.
- Khan, Uzma and Ravi Dhar. 2006. "Licensing effect in consumer choice." *Journal of Marketing Research* 43 (2):259–266.
- List, John A and David Lucking-Reiley. 2002. "The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign." *Journal of political Economy* 110 (1):215–233.
- Madrian, Brigitte C and Dennis F Shea. 2001. "THE POWER OF SUGGESTION: INERTIA IN 401 (k) PARTICIPATION AND SAVINGS BEHAVIOR." *Quarterly Journal of Economics* 116 (4).
- Monin, Benoit and Dale T Miller. 2001. "Moral credentials and the expression of prejudice." *Journal of personality and social psychology* 81 (1):33.
- Mullen, Elizabeth and Benoît Monin. 2016. "Consistency versus licensing effects of past moral behavior." *Annual review of psychology* 67.
- Nisan, Mordecai and Gaby Horenczyk. 1990. "Moral balance: The effect of prior behaviour on decision in moral conflict." *British journal of social psychology* 29 (1):29–42.
- Nisan, Mordechai. 1985. "Limited morality: A concept and its educational implications." *Moral education: Theory and application* :403–420.
- Ploner, Matteo and Tobias Regner. 2013. "Self-image and moral balancing: An experimental analysis." *Journal of Economic Behavior & Organization* 93:374–383.
- Pollak, Robert A. 1970. "Habit formation and dynamic demand functions." *Journal of political Economy* 78 (4, Part 1):745–763.
- Sachdeva, Sonya, Rumen Iliev, and Douglas L Medin. 2009. "Sinning saints and saintly sinners the paradox of moral self-regulation." *Psychological science* 20 (4):523–528.
- Shang, Jen and Rachel Croson. 2009. "A field experiment in charitable contribution: The impact of social information on the voluntary provision of public goods." *The Economic Journal* 119 (540):1422–1439.
- Strahilevitz, Michal A and George Loewenstein. 1998. "The effect of ownership history on the valuation of objects." *Journal of consumer research* 25 (3):276–289.
- Sunstein, Cass and Richard Thaler. 2008. "Nudge." *The politics of libertarian paternalism*. New Haven .


- Thaler, Richard H and Cass R Sunstein. 2003. "Libertarian paternalism." *The American Economic Review* 93 (2):175–179.
- Yörük, Barış K. 2013. "The impact of charitable subsidies on religious giving and attendance: Evidence from panel data." *Review of Economics and Statistics* 95 (5):1708–1721.
- . 2014. "Does giving to charity lead to better health? Evidence from tax subsidies for charitable giving." *Journal of Economic Psychology* 45:71–83.
- Zanna, Mark P. 1972. "Inference of belief from rejection of an alternative action." *Representative Research in Social Psychology* .

Appendix A Appendix

This section is meant for online publication only.

FIGURE A1: DONATION EXPERIMENT SCREENSHOTS

Thank you again for participating.



In addition to the participation payment of \$1 that you will receive, you have also earned another \$1.

The extra \$1 will be added to your final payment, so you will receive a total of \$2 which includes \$1 for participating in this study plus this extra \$1 bonus.

Your Extra Bonus Earnings: \$1

(A) ROUND 1, CASH ENDOWMENT

Thank you again for participating.



In addition to the participation payment of \$1 that you will receive, we will also make a donation on your behalf to a charity called CARE.

Your Donation Amount: \$1.50

This \$1.50 will be added to all of the donations of every participant to make a single payment to CARE. If you wish to receive confirmation of this donation, you will have the opportunity to indicate this in the survey at the end of this study.

(B) ROUND 1, CHARITY ENDOWMENT



You now have the option to give away your extra \$1 bonus. If you decide to give away your bonus, we will instead make a \$1.50 donation to CARE on your behalf. In that case, you will have a \$1 bonus from the participation fee and a \$1.50 donation to CARE. We will add your \$1.50 donation to all of the donations of other participants to make a single payment to CARE. If you wish to receive confirmation of this donation, you will have the opportunity to indicate this at the end of this study.

- I do not want you to take \$1 away from my bonus.
- I want you to take away my entire \$1 bonus in order to donate \$1.50 to CARE.

(C) SWAP CASH FOR DONATION

You now have the option to take away your \$1.50 donation to CARE. If you decide you do not want us to make a \$1.50 donation to CARE on your behalf, we will instead add \$1 to your bonus that you will get at the end of the survey. In that case, you will have a total of \$2 for your bonus consisting of \$1 of bonus for participation and \$1 for this decision, and we will not make any donation to CARE on your behalf.

- I do not want you to take away my \$1.50 donation to CARE.
- I want you to take away my entire \$1.50 donation to CARE in order to add \$1 to my bonus payment.

(D) SWAP DONATION FOR CASH

TABLE A1: ROUND 2: MULTIPLE PRICE LIST FOR DONATION EXPERIMENT

Option 1:	Add \$1.00 to your bonus and Donate \$0 to Save the Children.
Option 2:	Add \$.90 to your bonus and Donate \$.20 to Save the Children.
Option 3:	Add \$.80 to your bonus and Donate \$.40 to Save the Children.
Option 4:	Add \$.70 to your bonus and Donate \$.60 to Save the Children.
Option 5:	Add \$.60 to your bonus and Donate \$.80 to Save the Children.
Option 6:	Add \$.50 to your bonus and Donate \$1.00 to Save the Children.
Option 7:	Add \$.40 to your bonus and Donate \$1.20 to Save the Children.
Option 8:	Add \$.30 to your bonus and Donate \$1.40 to Save the Children.
Option 9:	Add \$.20 to your bonus and Donate \$1.60 to Save the Children.
Option 10:	Add \$.10 to your bonus and Donate \$1.80 to Save the Children.
Option 11:	Add \$0 to your bonus and Donate \$2.00 to Save the Children.

TABLE A2: LOCAL AVERAGE TREATMENT EFFECTS: ROUND 2 DONATION RATES & AMOUNTS

<i>Panel A: Moral Consistency</i>			
	Propensity to Donate	Donation Amount	
	(1)	(2)	(3)
$\widehat{A}_{i,1} \times StrongValue$	-0.18 (0.39)	0.04 (0.53)	-0.87 (2.25)
$\widehat{A}_{i,1} \times WeakValue$	1.17* (0.63)	1.50 (0.96)	7.89 (4.87)
Strong Value	0.42*** (0.15)	0.43* (0.22)	2.46** (1.05)
Constant	0.07 (0.08)	0.1 (0.12)	-2.68*** (0.78)
Observations	245	245	245
R^2	.	0.002	.
χ^2 test			
$\widehat{A}_{i,1} \times Strong = \widehat{A}_{i,1} \times Weak$	3.30*	1.75	2.66
<i>Panel B: Immoral Consistency</i>			
	Propensity to Keep	Keep Amount	
	(1)	(2)	(3)
$\widehat{c}_{i,1} \times StrongValue$	0.17 (0.24)	0.02 (0.27)	-0.44 (1.12)
$\widehat{c}_{i,1} \times WeakValue$	0.46 (0.46)	0.75 (0.48)	3.95 (2.44)
Strong Value	0.23 (0.45)	0.52 (0.47)	3.15 (2.27)
Constant	0.51 (0.42)	0.2 (0.43)	-1.61 (2.12)
Observations	245	245	245
R^2	0.03	0.002	.
χ^2 test			
$\widehat{c}_{i,1} \times Strong = \widehat{c}_{i,1} \times Weak$.32	1.75	2.66

Columns (1)-(4), OLS regression estimates. Column (5) & (6) Tobit regression estimates. Robust standard errors in parentheses and *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.