

In search of the limits of applying reciprocity in the field: Evidence from two large field experiments¹

Michael Sanders

Behavioural Insights Team and Harvard Kennedy School of Government

Abstract:

Experiments in both the lab and the field have gone some distance to proving that people are reciprocal agents, returning one good deed with another, even when it is disproportionately costly to do so. This finding is increasingly being applied to improve relationships between government and citizens, and to facilitate the smoother running of public services. We investigate the extent to which interventions intended to induce reciprocity work in the longer term, by asking participants to donate a comparatively large sum of money (£515 on average), to charity, after receiving a small gift (a packet of sweets). We find that people given a small packet of sweets by a volunteer are significantly more likely to donate to charity than those given a flier, and those not greeted by a volunteer. We find that our initial results replicate successfully on participants who were not treated in the initial experiment. Our principal contribution, however is to show that these effects fall by more than half the second time a participant receives a gift, although it remains of net financial benefit for the charity.

¹ This research is funded by the ESRC. I am grateful to Sabira Rouf and Amy Parker for making this experiment possible, to Olivia Maynard and Rachael Hogg for research assistance, to Michael Norton, Elizabeth Linos, Elspeth Kirkman, Syon Bhanot and Aisling Ni Chonaire and Ariella Kristal for comments on an earlier draft, and to seminar participants at the University of Bristol, Yale and at Harvard's Behavioural Insights Group for helpful feedback. I am also grateful to my PhD advisor, Sarah Smith, and to my PhD examiners Michael Price and Simon Burgess for their insights.

In search of the limits of applying reciprocity in the field: Evidence from two large field experiments

Scripture tells us that we should do unto others as we would have them do unto us. Not only this, but the social contract on which well functioning societies depend consists in part of myriad small kind acts being reciprocated over and over again. Single instances of these kinds of ‘gift exchanges’ have therefore been widely studied, but the evidence on their longer term impacts is rarer. A standard economic consideration of relationships based on trade suggest that one party – these seller – would only provide their goods or services if another – the buyer – has agreed to pay for them. Giving potential buyers a gift, and then hoping they pay for it afterwards, is counter-intuitive. Conditional transfers – where I agree to give you something if you give me something (often money) back, are the norm, while unconditional transfers (in which I give you something and hope that you give me something back), are rarer. Yet, this is the basis of reciprocal relationships are commonly seen in society. One case in which we might imagine that an agency is attempting to induce reciprocity is the phenomenon of gift-giving by charities to potential donors – for example sending potential donors a branded pen along with solicitation material. The available field experimental evidence on this supports the idea that reciprocity both exists and can be induced, but is contradictory as to whether or not it is a profitable tool for charities aiming to raise funds, and how it fares compared to more traditional conditional transfers. (Alpizar et al, 2012, Falk, 2008, Landry et al, 2014).

In spite of this, explicit field-experimental evidence on reciprocity with high stakes is comparatively scarce. Moreover, there is a general shortage of longer-term evidence on the effects of ‘behavioural’ interventions when they are applied to field contexts (Frey & Rogers, 2014). Given that reciprocal relationships are inherently based on goodwill, we might expect them to be particularly fragile in repetition, as gift receivers begin to feel as though they are being tricked. In this paper, we investigate these questions through two sequential field experiments carried out on the same population, attempting to encourage charitable donations. These experiments allow us to investigate participants’ pro-social behaviour, how they respond to receiving a small gift when asked to donate, and how this response changes when the intervention is applied repeatedly.

The structure of this paper is as follows. In the next section I briefly review some of the relevant literature on reciprocity. This is followed by the design and results of the first experiment in section 3. Section 4 describes the second experiment, while section 5 offers conclusions.

Previous studies

This paper extends the literature in two primary ways. First, it increases the body of field evidence on the existence of gift exchange and reciprocity, and secondly, it offers a comparatively rare field test of the effect of interventions when they are repeated.

The first area, gift exchange, is widely studied. In an environment where players are engaged in strategic play, it makes sense for rational economic agents to cooperate if doing so increases the likelihood of others doing kindly unto us at a later date. Lab studies find that in repeated play, cooperation in a public goods game is higher than if the game is not repeated (Dawes & Thaler (1988)). Conversely, where an individual is not identifiable (i.e. cannot be separated from the crowd), their cooperative behaviour falls significantly (Andreoni & Petrie (2004) Charness & Gneezy (2008)). Strategic cooperation also falls as the end of a finitely repeated game approaches (Fehr and Gächter (2000)). Taken together, our experiments cover a special case of this phenomenon, where participants are playing a repeated game, but where they, unlike the fundraisers, are unaware of this fact when they make their initial donation decision. This allows us to identify the effect of repetition in the absence of future expectations.

Theories of reciprocity distinguish between reciprocity in strategic play and the same action where there is no strategic relationship between players (Dufwenberg & Kirchsteiger (2004), Falk and Fischbacher (2006), Rabin (1993)). Even when there is nothing to be gained, reciprocity appears to persist in both the lab (Fehr, Gächter & Kirchsteiger (1996), Charness (2006) Charness and Haruvy (2002), Falk and Gächter (2003)) and the field (Frey & Meier (2004), Hennig-Schmidt, Sadrieh & Rockenbach (2010), Green, Gerber & Larimer (2008)), although the existence of true reciprocity remains somewhat controversial (Shaw and Santos (2012)).

Reciprocal action may also be a result of conditional cooperation strategies adopted by players. Theories of conditional cooperation (Dufwenberg et al (2001), Gächter (2007), Hugh-Jones & Reinstein (2012)), model people as being of three simple types – free-riders, co-operators and conditional co-operators. Free-riders will never contribute to a public good (donate to charity) and co-operators will always contribute. Conditional co-operators will contribute to the public good if others do the same.

Two field experiments that we are aware of test reciprocity in the field of charitable donations. In the first (Falk (2007)), participants are drawn from a ‘warm list’ of potential donors to a Swiss Children’s charity, who are sent either no gift, a small gift (one postcard) or a large gift (four postcards) and asked to donate to the charity. The large gift is found to significantly increase the proportion of participants making a donation. In the second (Alpizar et al (2008)) visitors to a Costa Rican national park are given a fridge magnet prior to being asked to give. Once again, the gift significantly increases the proportion of people donating, but decreases the average donation made. In the first study, the money raised through donations is greater than the cost of providing the gifts, while in the second it is not. Although the gifts cost about the same amount, and the proportion of people induced by the treatment is similar in both studies, the value of donations in the second was lower. This suggests that reciprocal action may be relatively price insensitive. These experiments, and others testing reciprocity, invariably concern themselves with one-shot encounters in which an individual is asked once to reciprocate and then not again.

Although the return on investment from inducing reciprocity may be positive in the first instance, people’s reciprocity may decay over time, as has been found of their tendency to “Pay It Forward” (Christensen (2014), Gray, Ward and Norton (forthcoming)). Specifically it is unclear whether if people reciprocate once, whether they would do so again in the future. The extent to which people habituate to interventions intended to alter their behaviours is both important for policymakers and as yet under researched (Frey and Rogers (2014)). Large scale randomised controlled trials in the UK (Hallsworth et al (2014)) have shown that participants who are late in their tax payments are more likely to pay if they are provided with social-norm messages - information that “Nine out of ten people have already paid their tax”. These changes have been estimated to increase the amount of tax being paid on time by £200 million per year in the UK (Daily Telegraph (2012))). Social norm messages have also been used in other areas (John et al (2014)), and are being used ever more widely. It is an open question, however, whether repeated exposure to social information reduce the effect of these messages. While we might expect more traditional government tools, such as incentives (through taxation and subsidies), to have a lasting effect, tools which make use of behavioural economics and psychology may have more transitive impacts.

An individual or organisation giving a gift for a second may have a stronger or a weaker effect on the donation behavior of others than doing so the first time. If gift giving is a means of sending a signal to the receiver about the giver’s type – that they are a nice person – doing so repeatedly may strengthen that signal, as the credibility of the

signal is amplified by the increased cost. However, if that signal is cheap, as gift giving appears to be, relative to the benefits in our experiment, its repetition could convince recipients that the giver is acting strategically in order to send an untrue signal of their type. We might expect gift exchange to be a powerful motivator of donations if participants are in some sense myopic. In game theory terms, if a participant is (or believes themselves to be) playing a one-shot game, and observes a generous act by another player, this will be taken as a credible signal of their kindness, to be reciprocated if the first player is a conditional cooperator. If a second, unexpected round of the same game is played subsequently, the first player might update their beliefs about whether or not they are playing a repeated game, and hence reinterpret the actions of the gift-giving player, which in an indefinitely repeated game might be taken as ‘cheap talk’ – a less credible signal of that player’s kindness. Interestingly, this hypothesis is easier to test in the field than the lab, as deception of participants is permitted, and participants may plausibly believe that a given interaction (in our case with a specific fundraising campaign), is a ‘one off’.

I now proceed to describe the experiments.

Experimental Setup

The workplace is attractive as a context for field experiments, for several reasons. During the working day it functions as a mostly closed environment – participants will usually have defined job role and not deviate particularly from it. Social communication is often limited, either by rule or necessity. In large firms, offices may be segregated geographically and may have little communication between them. Formal hierarchies are comparatively easy to observe, as are people’s approximate networks, based on who they work with. Importantly, we are also able to identify people’s approximate relative levels of income, according to the job title. Finally, data gathering on outcomes is possible without alerting participants to the nature of the experiment.

I report the results of two field experiments, which occur sequentially on much the same sample. Participants in both experiments are employees of a large investment bank in the UK. Participants are unaware that they are taking part in an experiment, and therefore this is a “natural field experiment” according to Harrison and List’s (2004) taxonomy. An investment bank is a particularly useful setting, as many of the regulations intended to prevent insider trading or similar activity, have the consequence of limiting communication between different parts of the bank that serve different functions.

Experiment 1

Our first experiment takes place as part of the bank's "One Day" fundraising campaign, in which employees are asked to donate a day's salary to two charities voted on by them earlier in the year. Our sample is unrepresentative of the general population, with the average gross daily salary of donors (1/250 of their annual salary) being approximately £515.50. Donations can be made either by clicking the link at the bottom of the CEO's email, by visiting the campaign's website, or by swiping their bank ID card (required to enter bank buildings) on portable scanners located around the building.

Participants are divided into 67 business units, many of which have similar functions, but which are geographically disparate, and anecdote suggests do not communicate with each-other. For logistical reasons, and to avoid contamination within clusters, our treatments in both experiments are assigned randomly at the business unit level, although donation decisions are observed at the individual level.

The experiment was carried out over a single day in August 2012. This experiment was conducted on 6175 individuals, who are employees of a large investment bank in London. Ancillary and secretarial staff of the bank are excluded from the experiment, and so participants' job titles within the bank (in ascending order of pay and status) are analysts, associates, assistant vice presidents, vice presidents, directors and managing directors. Details of starting salaries for these bands are in the table below. However, I note that these figures are illustrative only, as the distribution of salaries within each title is incredibly heterogeneous, and that the top of each pay band considerably overlaps the bottom of the band above. As such, these figures are interpreted primarily ordinally- such that an associate will on average earn more than an analyst but less than an assistant vice president, and so on. Participants are included in our analysis if they are present at work on the day of the experiment, ascertained by whether or not they swipe their ID card to enter the building on that day. Within this data, participants are identified by their job title and the business unit in which they work, and they are linked to treatment conditions by their business unit. Where participants' business unit is identified as being outside of our sample frame (for example, when it is in New York), participants are dropped from analysis. As none of these individuals donate, this does not materially impact our findings.

Table 1: Job Titles and Salaries

Job Title	Mean Starting Salary
Analyst	£45,000
Associate	£62,000
Assistant Vice President	£69,000
Vice President	£95,000
Director	£122,000
Managing Director	£265,000

Interventions – Experiment 1

This experiment tests whether participants reciprocate a small gift when their choice of donation amount is constrained to large units of their salary. Donations made were matched by the bank and split between two charities – Meningitis Research UK and Help a Capital Child. Business units were assigned to one of three treatments – control, volunteers, and sweets (reciprocity).

Our control condition describes the ‘campaign as normal’. All participants in this condition receive an email from the bank’s CEO, asking them to donate a day’s salary to the two charities. Emails were sent automatically from the CEO’s company email address, at 8 o’clock on the morning of the campaign. The email included the URL of the donation page, and a link to click to donate (if a participant clicked this link, a donation was made without any further action being needed from the participant). There was also a link to say ‘not this time, thank you’. The email emphasised the work that could be done with donations “funding vital research” and “supporting 20,000 disadvantaged children”. Participants could donate responding to the email, by swiping their ID card on scanners located around the building, by clicking a link in an email sent to them by the bank’s CEO, or by visiting the campaign’s website and entering their email address. A day’s salary (defined as 1/260th of the employee’s pre-bonus salary), was then deducted from the subsequent month’s pay automatically.

The fundraising campaign was also advertised around all of the bank’s UK offices with posters displayed in common spaces such as coffee areas, bathrooms and lifts. The branding of the campaign emphasised the ease of giving through this campaign, compared with other ways of fundraising, such as climbing mountains, running

marathons, or baking cakes (see appendices for details). Considerable effort was exerted by the bank's CSR department to reduce the complexity of donating in this campaign, and this marketing decision was based on focus groups of bank employees, which suggested that employees valued this ease.

Participants in the *Sweets* treatment were greeted by volunteers on entering their business unit and offered a small packet of sweets. Both the clothing (T-shirts), of volunteers, and the packet of sweets, were branded with the "Take the Easy Way" branding of the campaign. A picture of the packet is in figure 2, and the sweets themselves were Skittles, a popular brand of sweets in the UK. The cost of each packet of sweets was £1.49 (\$2.25). We note however that transactional value of the sweets to the participants may be lower. The sweets were of a readily available make, and a packet of equivalent size could be bought easily for less than 50 pence (75 cents) – the packaging and branding was the majority of the cost to the experiment. A picture of a volunteer in this experiment can be found in figure 3.

As in Alpizar et al (2008) hand-delivery of sweets by volunteers was the only way of ensuring that participants were all offered sweets and understood the context in which they were given (and were given the opportunity to ask questions). However, it is possible that the face-to-face interaction with a volunteer, rather than a sense of reciprocating the sweets, might be driving any effect observed. This is addressed by our third treatment.

Participants assigned to the *Volunteer* condition were greeted on arrival at work by a volunteer, dressed in the same branded t-shirt, and were given a scaled down version of the posters displayed in other conditions. A flier is shown in figure 4. Volunteers in both conditions were given the same training and information about the campaign, which is described in the supplementary materials.

Volunteers were provided by the charities supported by the campaign, and, where required, by the research team.



Figure 1: Packet of Sweets, Experiment 1



Figure 2: Volunteer, Experiment 1



Figure 3: Flier, Experiment 1

Results – Experiment 1

I compare the three conditions using simple regression analysis, with standard errors clustered at the level of the business unit. 6.1% of participants in the control group make a donation. I find no significant difference between behaviour in the control group and those in the volunteer treatment (6.1 and 6.15%, $p=0.61$), indicating that simply being asked to donate by a human being has no influence on donations in this case. By comparison, we find that the gift of a small packet of sweets has a large and statistically significant effect (6.1 and 13.54%, $p<0.01$). This effect is robust to different specifications of the regression model. To protect individual employees' anonymity, data on

participants are minimal, except for their rank in the business. The effect of reciprocity is largest among low ranking employees, and declines as participants become more senior. As randomization was conducted at level of the business unit, there is a higher risk of sample imbalance in this experiment than were the randomization conducted at the individual level. Analysis of the sample's balance shows that the volunteers treatment is imbalanced with respect to the Analysts ($p=0.016$), while the Sweets treatment is imbalanced with respect to both the Analysts ($p=0.033$) and Managing Directors ($p=0.012$). As such, we conduct analysis controlling for participant rank. Doing so reduces the estimated effect size, but it remains significantly different to the control and volunteer groups (4.6% $p<0.01$), and this is not significantly different to the non-adjusted effect size ($p=0.08$).

Table 2: Impact of Reciprocity Treatment on Likelihood of Donating, Experiment 1(OLS)

	(1) (All)	(2) (Analysts)	(3) (Associates)	(4) (Ass. Vice Presidents)	(5) (Vice Presidents)	(6) (Directors)	(7) (Managing Directors)
Volunteers	0.00541 (0.00790)	0.0658 (0.108)	0.0170 (0.0183)	0.000997 (0.0126)	0.0124 (0.0134)	-0.0391 (0.0238)	0.0503 (0.0695)
Sweets	0.0631*** (0.0106)	0.243* (0.114)	0.0521 (0.0304)	0.0196 (0.0269)	0.0238 (0.0179)	0.0713** (0.0261)	-0.0874 (0.0675)
Constant	0.0614* (0.00734)	0.131 (0.115)	0.00644 (0.0136)	0.0304* (0.0138)	0.0225 (0.0128)	0.0445* (0.0204)	0.0862 (0.0467)
<i>N</i>	6175	87	639	1020	1949	1305	500

Discussion – Experiment 1

The findings of this experiment suggest two things – first, that reciprocity of a small gift appears to be relatively insensitive to the amount of money asked for, in line with theories of conditional cooperation. Second, the reduced effect at higher levels of the firm's hierarchy suggests either that participants at this level are less likely to be conditional co-operators, or that the size of the gift relative to a person's income does matter to some extent. Given findings that conditional co-operators or 'matchers' are more likely to succeed in a firm than to be mid-ranking (Grant (2013)), the latter interpretation is more intuitive.

Cost Benefit Analysis – Experiment 1

Although we cannot observe the value of individual donations, it is possible to make conservative estimates of the cost/benefit ratio for this treatment.

The point estimate for our reciprocity treatment is 0.063. Hence, one donation is induced for every 15.8 packets of sweets distributed. Each packet of sweets costs £1.50 so the cost for every induced donation is £23.77. If all marginal donors were Analysts (the lowest grade) paid the starting salary for that grade of £45,000, each induced donation would be worth £173.07 (donations are of salary divided by 260), plus a 100% match from the bank, making the value to the charity £346.14. Hence, the minimum benefit/cost ratio of this treatment is 14.5:1. I therefore conclude that in this context, unlike that in Alpizar et al (2008), the benefit to the charity outweighs the cost of the gift.

Experiment 2

To answer the question of whether repetitive gift giving has continued effectiveness, as well as to try and replicate our first experiment, a second experiment was carried out over a single day in August 2013. Participants were 6833 employees of the same investment bank, subject to the same restrictions as experiment 1. The basic structure of the fundraising campaign was the same. Donations made were matched by the bank and split between two different charities – Elephant Family and Helen Bamber Foundation, which were selected by employees of the firm.

Interventions – Experiment 2

The same business units as in experiment were allocated to two conditions: control and sweets. Units allocated to the sweets condition in year 1 remained allocated to the sweets condition, and all other units were randomly re-allocated. Denoting a unit that received sweets as S and one that did not N, our data contain units which can be denoted as NN, NS, SS, but not SN. Although this prevents us from identifying the effect of withdrawing sweets from someone who previously received them, we are able to identify the effect of participants becoming accustomed to gifts of this type.

Participants in the *Control* condition were sent a generic (not personalized) email from the CEO at the beginning of the campaign as they had been previously. Posters and fliers were distributed around the offices of the bank during the week leading up to the campaign, and were located in shared areas – bathrooms, coffee shops and lifts. Senior managers in the bank were sent emails in the campaign inviting them to encourage their colleagues to donate. As

well as the charities, the branding of the campaign changes substantively between years, as can be seen in the figures.

Volunteers greeted participants in our two *Sweets* conditions on arrival in their offices. These volunteers wore T-shirts branded with the campaign's brand for 2013 (shown in figure 5). Volunteers offered participants who passed a small packet of sweets, shown in figure 6, from a bucket of identical sweets (figure 7). Once again, the cost of sweets to the campaign was £1.49 (\$2.25) per packet, and the value of the sweets in the packet was roughly 50 pence (75 cents). Volunteers were given information about the charities, and basic information on the fundraising campaign, to answer any questions the employees might have. Among employees who received sweets there were two groups – “*Sweets Once*” participants, whose business unit did not receive sweets in the 2012 campaign, but did in the 2013 campaign, and “*Sweets Twice*” participants, whose business unit received sweets in both campaigns.



Figure 4: Packet of Sweets, Experiment 2



Figure 5: Volunteer, experiment 2



Figure 6: Bucket of sweets, experiment 2

Results – Experiment 2

Our analytical strategy for this experiment is the same as for the first. A linear regression is run, with standard errors clustered at the level of the business unit, estimating the proportion of participants making a donation of a day's salary by the two sweets groups compared with the group who received no sweets. 4.4% of participants in the control group make a donation. Our first result is that the effect of reciprocity is replicated, with 6.6% more participants donating in the *Sweets Once* group than the control group ($p < 0.001$). This result is not significantly different to the results of our first experiment ($p = 0.08$), although it is larger in absolute value.

We find that participants in the *Sweets Twice* group are 2.6% more likely to make a donation than participants in the control group ($p = 0.026$). This group is also significantly *less* likely to make a donation than participants in the "sweets once" group ($p < 0.001$), and significantly less likely to make a donation than members of the same group had been on the first occasion when they received sweets ($p < 0.01$).

I also conduct similar analysis to the first experiment, splitting the results of our experiment by the rank of participants. Here we find that managing directors are insensitive to reciprocity treatments, consistent with our previous findings, and that directors who received sweets earlier were similarly unresponsive. We find that for all ranks except analysts (the most junior rank), participants receiving sweets for the first time were significantly more likely to donate than participants who had received them previously.

Table 3: Effects of Reciprocity on Probability of Donating, Experiment 2 (OLS)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	(All)	(Analysts)	(Associates)	(Ass. Vice Presidents)	(Vice Presidents)	(Directors)	(Managing Directors)
Sweets Once	0.066*** (0.016)	0.077 (0.080)	0.161* (0.071)	0.016 (0.036)	0.126* (0.059)	0.047 (0.057)	-0.084 (0.066)
Sweets Twice	0.027** (0.012)	-0.022 (0.086)	0.108* (0.044)	0.040 (0.025)	0.031 (0.032)	-0.009 (0.039)	-0.096 (0.089)
Constant	0.031*** (0.006)	0.016 (0.058)	0.058 (0.035)	0.007 (0.010)	0.010 (0.018)	0.022 (0.031)	0.153** (0.052)
Observations	6833	358	1066	1627	3027	1968	909

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, Standard Errors Clusters at the Business Unit Level

Discussion – Experiment 2

I argue that this result shows that attempting to induce reciprocity continues to work on repeated attempts, but that the effect is significantly and substantially reduced. This occurs even when the charities being fundraised for are substantially different, as is the branding of the campaign, and a year has passed in the intervening time. I further argue that this finding is in line with theories of reciprocity, such as that put forward by Falk & Fischbacher (2006), in which the receiver's perception of the giver's motives determines whether or not the receiver reciprocates a gift. Participants in this paradigm might respond to a gift once, but do not do so a second time, when they might begin to see gift giving as a strategic (rather than altruistic) act. We can think about this in the context of a repeated game – participants in the first experiment were unaware that there would be a repeated interaction (that the second experiment would take place), and therefore could treat the small gift as a credible signal of altruism from the charities/bank. When the second experiment takes place, the repeated nature of the interaction is more starkly revealed, and participants may update their beliefs about the credibility of this signal. An alternative explanation may be that participants become accustomed to receiving sweets in exchange for donations quickly, and so their power as incentives are reduced. In future experiments, I would hope to disentangle these potential causes.

Cost Benefit Analysis:

In Experiment 2, the average value of donations remains the same according to the information available from the firm, and the cost of the sweets remains unchanged. The point estimate for the reciprocity is slightly higher (0.066), and so the cost for every induced donation is 22.72. Following the same conservative estimate, the minimum benefit/cost ratio for the *Sweets Once* condition is 15.2:1. Among participants receiving sweets for a second time, receiving a gift is significantly less effective at inducing donations, and so the cost per donation induced rises to

£38.63, giving a benefit/cost ratio of 5.9:1 – still a highly profitable undertaking for the charity, but substantially less than in the first year.

Conclusions

I draw two main conclusions based on our findings from these experiments. Firstly, this paper contributes to the literature that suggests a large impact of reciprocity on human behaviour, and finds that people's desire to reciprocate may be relatively price insensitive, given the amount of money participants are asked to donate. Although the participants are unusually high earning, the amount of money they are being asked to donate is non-trivial (5% of their monthly salary). Secondly, in the process of replicating this finding, I discover that the effects of reciprocity do not appear fully to bear repetition within individuals, with the effect of a small gift on the proportion of employees donating falling by roughly half. The underlying causal mechanism is not yet fully , but I postulate that this shows further evidence that motive matters when attempting to induce participants to cooperate, since people may view a second attempt at using the same intervention as more explicitly manipulative. In addition, I am able to contribute a replication of our initial experiment on the same population, confirming the prior results. Similarly to prior research, we find that charitable donations are most common among the most and least highly ranked participants the firm. Reciprocal behaviour, however, is less common among more senior employees of the firm than the more junior ones. It is not possible to say the extent to which this is a result of selection (more reciprocal people failing to get promoted), or the rising discrepancy between the value of the gift and the value of the requested donation. Cost benefit analysis reveals that in our experiments both the initial giving of sweets and its repetition produce profit to the charitable causes, but that this profitability declines rapidly, and may disappear on further repetitions, or if donations are likely to be smaller than those we observe in this experiment. When considering these findings from a broader policy perspective, it suggests that interventions taken from behavioural science and applied to policy, while initially effective, may have smaller effects in the longer term. Either as the public become suspicious of them, or as people adapt their cognitive shortcuts to reflect changes in the environment brought about by policymakers. Empirical research on the longer term effects of behavioural science interventions would clearly be beneficial, as would a detailed theoretical consideration of the interplay between policymakers' and other agencies' use of behavioural biases and the public's responses to them. Further studies should look at different stakes and frequencies of intervention delivery, as well as allowing for a greater level of confidence in group composition stability between time periods.

References:

- Alpizar, F., Carlsson, F., & Johansson-Stenman, O. (2008). Anonymity, reciprocity, and conformity: Evidence from voluntary contributions to a national park in Costa Rica. *Journal of Public Economics*, 92(5), 1047-1060.
- Andreoni, J., & Petrie, R. (2004). Public goods experiments without confidentiality: a glimpse into fund-raising. *Journal of public Economics*, 88(7), 1605-1623.
- Charness, G., & Gneezy, U. (2008). What's in a name? Anonymity and social distance in dictator and ultimatum games. *Journal of Economic Behavior & Organization*, 68(1), 29-35.
- Christensen, R. K. (2014). Concepts to Advance Public Service Contribution? Equity, Motivation, and Altruism. *Journal of Public Administration Research and Theory*, 24(4), 1076-1079.
- Dufwenberg, M., Gneezy, U., Güth, W., & Van Damme, E. (2001). Direct vs indirect reciprocity: an experiment. *Homo Oecon*, 18, 19-30.
- Dufwenberg, M., & Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and economic behavior*, 47(2), 268-298.
- Falk, A. (2007). Gift exchange in the field. *Econometrica*, 75(5), 1501-1511.
- Falk, A., & Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, 54(2), 293-315.
- Fehr, E., & Gächter, S. (2000). Fairness and retaliation: The economics of reciprocity. *The journal of economic perspectives*, 159-181.
- Frey, E., & Rogers, T. (2014). Persistence How Treatment Effects Persist After Interventions Stop. *Policy Insights from the Behavioral and Brain Sciences*, 1(1), 172-179.
- Gächter, S. (2007). *Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications*. na.
- Grant, A. (2013). *Give and take: A revolutionary approach to success*. Hachette UK.
- Gray, K., Ward, A. F., & Norton, M. I. (2014). Paying it forward: Generalized reciprocity and the limits of generosity. *Journal of Experimental Psychology: General*, 143(1), 247.
- Hallsworth, M., List, J. A., Metcalfe, R. D., & Vlaev, I. (2014). *The behavioralist as tax collector: Using natural field experiments to enhance tax compliance* (No. w20007). National Bureau of Economic Research.
- Harrison, G. W., & List, J. A. (2004). Field experiments. *Journal of Economic Literature*, 1009-1055.
- Harper, H, A Sallis and M Sanders (forthcoming): "Encouraging organ donation register sign-up using behaviourally informed messages – a randomized trial" *working paper available on request*
- Hugh-Jones, D., & Reinstein, D. (2012). Anonymous rituals. *Journal of Economic Behavior & Organization*, 81(2), 478-489.
- John, P., Sanders, M., & Wang, J. (2014). The Use of Descriptive Norms in Public Administration: A Panacea for Improving Citizen Behaviours?. *Available at SSRN*.

Kettle, Ruda, Hernandez & Sanders (2015): "A Field Experiment testing the effect of nudge-type messages on tax compliance in Guatemala" *Working Paper*

Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American economic review*, 1281-1302.

Sanders, M& E Kirkman, (2015): "I've saved you a place, good luck: personalization and reciprocity in attendance for high-value job coaching sessions" *Working Paper, available on request*

Sanders & Hallsworth (2015) "Applying Behavioural Economics in a Health Policy Context" in "*Behavioural Economics and Health*" ed. Christina Roberto