

DISCUSSION PAPER SERIES

IZA DP No. 10736

**The Sequencing of Gift Exchange:
A Field Trial**

Jeffrey Carpenter

APRIL 2017

DISCUSSION PAPER SERIES

IZA DP No. 10736

The Sequencing of Gift Exchange: A Field Trial

Jeffrey Carpenter
Middlebury College and IZA

APRIL 2017

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

The Sequencing of Gift Exchange: A Field Trial*

There is now an extensive literature on “gift exchange” showing that when principals and agents can trade “gifts” (rewards that should not emerge in a competitive equilibrium), exchange becomes more efficient. However, it is not obvious how gift exchange should be organized if the principal’s goal is to increase the performance of a reciprocal agent. Specifically, who should make the first gift, the principal or the agent? Although both orderings, by themselves, have been hypothesized and examined in theory and experiments, the literature is largely silent on the comparison. I report the results of a field experiment that compares the principal-first and agent-first orderings to each other and a gift-less control. Consistent with the previous experimental literature, I find that principal-first, gifts do increase agent performance. Unlike the literature, however, I find that agent-first, gifts are also effective. Comparing the two, I see that the agent-first ordering works best, is clearly cheaper to implement and differences appear on both the extensive and intensive margins.

JEL Classification: C93, D03, D64, H41, L30, M30

Keywords: gift exchange, reciprocity, social norm, incentives, field experiment, charity, fundraising

Corresponding author:

Jeffrey Carpenter
Department of Economics
Middlebury College
Middlebury, VT 05753
USA
E-mail: jpc@middlebury.edu

* The paper is forthcoming in the Journal of Economic Behavior and Organization.

1 Introduction

Since Akerlof (1982) there has been considerable interest in how social norms of reciprocity can influence market outcomes. The reciprocal notion of “gift exchange” was first applied to labor markets (Akerlof and Yellen, 1986) but has since been linked to many other settings including product markets (Kahneman, Knetsch and Thaler, 1986; Anderson and Simester, 2008), informal financial markets (Lee and Persson, 2016) and philanthropic markets (Andreoni and Payne, 2013). One thing that remains, more or less, constant, however, is the sequencing of reciprocity: in most cases a principal offers a gift (some rent beyond the compensation necessary to clear the market) and an agent reciprocates by performing beyond what would be expected in equilibrium. While this sequencing often makes sense, for example bosses offer greater than market-clearing wages to a workers who respond by working hard, one can also find examples in which the reversed sequencing seems equally, if not more, natural. Consider workers providing effort beyond what their current compensation would justify being rewarded by the gift of a large (more than competitive) raise. Or, in a different context, what seems more natural: charities preemptively awarding gifts to potential donors or donors being given gifts in exchange for their contributions?

Interestingly, the sequencing is unlikely to matter in theory given the simplicity of the basic gift exchange formulation (as essentially a sequential prisoners’ dilemma - see Dufwenberg and Kirchsteiger, 2004), and, perhaps as a result, the empirical literature has yet to examine whether the principal’s gift coming first or second is more effective at motivating an agent’s reciprocal response. This paper describes a field experiment designed to examine whether the sequencing of gifts does affect the strength of agent reciprocity.

While, to my knowledge there is only one directly related paper in the economic literature (Landry et al., 2012 described in detail in the discussion), there are plenty of lab and field experiments in which the principal moves first. Charness and Kuhn (2011) provide a survey of such experiments formulated in the labor market context. Overall, there seems to be some debate as to whether the standard sequencing has a lasting effect on agent reciprocity. Gneezy and List (2006), for example, find in a real effort experiment that worker effort does increase in response to the gift of a wage rent by the principal, however the effect wanes after a few hours. Kube et al. (2012) also find limited support for gift exchange when the principal offers a monetary gift;

however, they find large and lasting effects when the gift is non-monetary or when some effort is devoted to the presentation of a monetary gift. Though not identical, Karlan’s (2005) study of the trust game in the microfinance context, provides some evidence of gift exchange in the field. The studies closest to this one, however, are the field experiments in the philanthropic context of Falk (2007) and Alpizar et al. (2008) wherein a charity first offers a gift in the hope of donor reciprocity. In both studies modest non-monetary gifts seem to be effective, though mostly on the extensive margin.

There are also a few studies in which the order of gifts is reversed. In the labor context, Fehr et al. (2007) allow principals to offer high “trust” wages before production occurs or “bonuses” after production and find that the principals drift towards the use of bonuses. Abeler et al. (2010) require workers to first choose a level of effort to be followed by the principal’s payment of a wage. The authors study how the forced equalization of wages across (possibly) heterogeneous workers can undermine reciprocal motives in a firm. Using a very similar design, Kleine and Kube (2015) show that self-assessments of performance can also be detrimental to gift exchange in a firm because workers tend to over-assess their performance and this makes managers distrustful and less generous with their subsequent wage gifts. In the philanthropic context, Mellström and Johannesson (2008) focus on whether financial rewards crowd out blood donations and find that offering payments after a donation reduces the supply of blood donated from women by almost half (while there was no effect on men). Considering cash donations, Newman and Shen (2012) study the more common ordering of gifts - charities giving gifts to people after they have made a donation. Although most participants expect such gifts to increase giving, the authors find exactly the opposite and attribute the result to the crowding out of intrinsic/altruistic motivation. In a related direct mail field experiment conducted to benefit a non-profit radio station, Chao (2016) documents crowding too: donations are less frequent (and a bit lower on average) when thank you gifts are offered, again most likely because of crowding. While reciprocity emerges to one degree or another in the labor context when the ordering of gifts is reversed, in the philanthropic context it seems that this reversal may be detrimental. However, because no direct comparison can be made it is impossible to know how the results would have differed had both sequences been examined.

To explore whether the sequencing of gift exchange matters, I report the results of a door-to-door fundraising experiment that benefited a local non-profit. Partici-

pants were randomized into three (non-monetary) gift treatments: one in which no gift was offered, one in which a gift was offered to every potential donor just before they were solicited and one where participants knew that only donors would receive a gift. Comparing the “No Gift” and the “Gift Before” conditions yields results that reproduce what others have found, to some extent: offering gifts before the solicitation yields a modest increase in the rate of giving, though not enough to significantly affect the overall average gift. At the same time, and in contrast to what others have found concerning conditional gifts, adding the “Gift After” data to the comparison indicates that this is the most effective sequencing of gift exchange: rates of giving and gifts are larger than both of the other conditions and because fewer gifts are actually made, the ordering is clearly the most cost effective method of initiating gift exchange. Lastly, all these results are robust to the inclusion of a number of important controls (e.g., being on the donor “warm list,” having benefited from the charity and demographics) that were also collected.

2 Methods

As a test of whether the sequencing of gift exchange affects agent reciprocity, together with my research solicitors, I designed and implemented a door-to-door fundraising experiment that took place in Addison County, Vermont during the Spring of 2016. The experiment consisted of randomizing individual participants (on the doorstep) into one of three conditions: A control in which no gift was offered, a treatment in which gifts were offered just before a solicitation was made and a second treatment that reversed the sequence such that gifts were given to donors after solicitations were made. The details of our procedures, the treatments and our participants are as follows.

All the proceeds from the fundraiser benefited a local poverty relief organization that was started in 1965 (to complement the Johnson administration’s war on poverty). The organization is well-known and highly regarded locally and has the vision, “that all people in Addison County have access to the tools and resources necessary to meet their own basic needs.” To implement this vision the charity provides food assistance (via a food bank) and other basic needs (such as heating fuel assistance), job training and counseling to individuals and families in the surrounding county. The organization runs a number of events each year (e.g., letter drives) and therefore our event would not have seemed out of the ordinary in the community.

To promote the fundraiser, in the weeks prior to and during the event we ran advertisements in the local newspaper and on a local radio station. We also hung posters at gathering places in the county (e.g., grocery stores and shop windows) and the charity publicized the event on its Facebook page. Lastly, when in the field the solicitors wore t-shirts bearing the name of the charity and carried charity-made identity cards to further advertise and assure participants that the fundraiser was legitimate.

At the beginning of April thirty research solicitors were trained. They were instructed to follow a standard protocol for every home visited and to use a script that varied only in the explanation of the treatments. Our control treatment implemented a standard voluntary contribution mechanism. Here, after briefly describing the charity, the solicitors asked for a donation. Specifically, they asked, “Would you like to make a donation to {Charity Name} today?” In the Gift Before treatment, just prior to asking for a donation the solicitors offered each participant a non-monetary gift - a silicon wristband embossed with the charity’s name and logo. In this treatment the solicitors stated, “Here is a {Charity Name} wristband for you. Would you like to make a donation to {Charity Name} today?” Lastly, in the Gift After treatment, the sequence was reversed and participants were told before being solicited that they would receive the gift if they made a donation. The solicitors said, “We’re offering {Charity Name} wristbands to people who donate. Would you like to make a donation to {Charity Name} today?”

During the end of April and the beginning of May, the research solicitors, in teams of two, approached more than 2000 homes in Addison County, Vermont and during 1107 of these home visits, someone answered the door. Considering there are just 37k people in the county living in 17k households (according to the Census Bureau estimates for 2015), our 1107 solicitations represents a relatively large sample. Prior to knocking on the door, the research solicitors randomized each home into one of the three conditions by rolling a die. This was done so that each solicitation team ran all three of the treatments and we can estimate treatment effects within solicitation team. When a door was answered, the research assistants followed the standardized script that varied only in the explanation of the gift exchange treatment as described above. Following the solicitation, the research assistants recorded whether or not a donation had been made and the amount of the donation. The solicitors accepted cash or checks made out directly to the beneficiary. The encounter ended after the solicitors asked the participant if he or she had ever benefited from the services provided by our charity.

Once off the stoop, the solicitors recorded the home’s address and the estimated the age and gender of the donation decision-maker. As another potentially important control, we made use of the charity’s administrative records to determine whether any of the households had previously given directly to the organization. We collected this data on previous donations to account for and estimate any “warm list” effects (à la Landry et al., 2010).

Fifty-four percent of our participants were women, 7% had previously taken advantage of our charity’s services, 13% were on the warm list (i.e., they had donated to the charity before) and the average age of the participants was estimated to be 53 years. Considering our randomization, the lowest p-value resulting from pairwise t-tests of the participant observables was 0.08 concerning the difference in proportions of participants who had benefited from the services provided by our charity in the Gift Before and Gift After treatments. In other words, our randomization seems to have worked well because we achieved balance on the observables that we collected.

3 Results

Combining the donations from all the experimental conditions, our fundraiser was relatively successful compared to both other experiments listed in the introduction and other events run by the non-profit. We gathered \$4400 in donations for the charity. The average donation, including zeros from the 60% of the participants who chose to not donate, was \$3.97. Considering just the positive donations, the average gift was \$9.82. Further, all our wristband gifts were accepted by the people solicited. The analysis begins with a discussion of our main result - the effect of the ordering of gift exchange on the amount donated - including all the solicitations and then we examine what drives our main result, participation (i.e., the extensive margin), increased giving (the intensive margin) or both?




Figure 1: Donation Behavior Differences by Treatment (bars represent the 95% confidence interval).

Our results are summarized in Figure 1. Focussing attention on all the solicitations (i.e., including those that did not yield a donation), Panel (a) of Figure 1 charts the mean donations in each condition, along with the 95% confidence intervals. The average amount received in the No Gift condition was just under three dollars, \$2.89 to be precise. The principal-first sequencing does improve on this. In the Gift Before treatment the average amount received was \$3.59, though the \$0.70 difference, while economically significant, is not statistically significant ($t = 1.41$, $p = 0.16$). The Gift After treatment, however, is a further improvement, one that does better than both the No Gift control and the Gift Before sequencing. Here the average solicitation yielded \$5.47, an amount that is significantly greater than both other conditions. The \$2.58 increment over the No Gift treatment is large and highly significant ($t = 3.05$, $p < 0.01$) and the \$1.88 “bump” above the Gift Before treatment is also significant using a simple t-test ($t = 1.97$, $p = 0.05$).

Our t-tests results are replicated in column one of Table 1 which reports regression results that include the full sample, other influences and examine the robustness of our summary tests. The table also includes estimates (near the bottom) of the difference between the point estimates from the two sequencing treatments. In column four we add solicitor fixed effects and the other variables we were able to gather and find their effects to be mostly orthogonal to the sequencing effects, as randomization to condition should assure. The one difference we see, however, is that the estimated effect of the

Gift Before treatment increases to \$1.02 which is now significant at the 10% level. The inclusion of the solicitor fixed effects drives this change. We find no significant effect of demographics: donations decrease in estimated age but the point estimate is tiny and though women are estimated to give \$0.50 more the effect is not significant ($p = 0.46$). We also do not find that participants who have previously benefitted from the services of our charity give an amount that is significantly different from other participants. However, we do find that previous donors (those on the warm list) give considerably more than others. Those participants gave \$3.93 more than the others ($p = 0.04$).

Our results suggest that the sequencing of gift exchange matters: conditional gifts seem to work better than unconditional ones. We now examine whether this difference in performance is driven by differences in participation or the amount given, conditional on participating. Considering the extensive margin first, Panel (b) of Figure 1 reveals that part of the reason why the Gift After treatment does well is because people are more likely than in the other conditions to make a donation. The participation rate was 48% in the Gift After treatment and only 41% and 34% in the Gift Before and No Gift treatments, respectively.

Table 1: Gift Exchange Treatment Effects

	Donation	Donate?	Don.>0	Donation	Donate?	Don.>0
Gift Given Before (I)	0.698 (0.502)	0.074** (0.036)	0.155 (1.041)	1.021* (0.534)	0.069* (0.036)	1.052 (1.182)
Gift Given After (I)	2.581*** (0.873)	0.140*** (0.035)	2.880 (1.751)	2.758*** (0.991)	0.135*** (0.034)	3.250* (1.853)
Estimated Age (#)				-0.005 (0.016)	-0.003*** (0.001)	0.067 (0.041)
Female (I)				0.502 (0.685)	0.036 (0.029)	-0.297 (1.770)
Benefited (I)				-0.060 (1.268)	0.092 (0.064)	-0.818 (1.925)
Warm List (I)				3.931** (1.954)	0.081* (0.046)	6.467* (3.455)
Constant	2.895*** (0.314)	0.336*** (0.023)	8.623*** (0.720)	2.775** (1.255)	0.527*** (0.076)	4.350 (2.700)
Est. (After-Before)	1.882** (0.903)	0.067* (0.038)	2.724 (1.765)	1.737** (0.882)	0.066* (0.037)	2.200 (1.611)
Solicitor Fixed Effects	No	No	No	Yes	Yes	Yes
Observations	1107	1107	448	1107	1107	448

Notes: Three dependent variables are considered: donation amounts, including zeros, whether or not a donation was made and the amount of just the positive donations. All results are from OLS with robust standard errors.

Returning to Table 1, results on the extensive margin are reported in columns two and five. Column two suggests that the 7.4 percentage point increase in participation accruing to the Gift Before sequencing is significant at the 5% level but that the participation increment is doubled compared to the No Gift treatment when the order

of gifts is reversed in the Gift After treatment. This 14 percentage point increase in participation is significant at better than the 1% level. Further, the 6.7 percentage point difference between the two gift treatment estimates is also marginally significant ($p = 0.08$). Promising the gift after performs better partially because it attracts more participation.

Adding controls to the participation regression in column five of Table 1 does little to affect the gift exchange treatment point estimates. We do see, however, that older people in our sample are significantly less likely to participate ($p < 0.01$) but the measured effect is not strong: a decade increase in age reduces the likelihood of giving by just 3 percentage points. Lastly, the effect of having a history of giving to the charity is strong on the extensive margin. We find that people on the warm list are 8 percentage points more likely to give ($p = 0.08$).

The differences that developed on the intensive margin are reported in Panel (c) of Figure 1. What is apparent in this panel is that the average positive gift does not differ much between the No Gift baseline and the Gift Before treatment, as is also common to the studies discussed above, but what is unique to our experiment is that the Gift After sequencing has an appreciable, if a bit noisy, effect on the intensive margin. Mean giving conditional on making a donation is \$8.62 in the No Gift baseline, \$8.78 in the Gift Before treatment and \$11.50 in the Gift After treatment. Using summary t-tests, the No Gift - Gift Before difference is far from significant ($p = 0.88$), but the large difference between the Gift After and both the No Gift and Gift Before conditions is much closer to significant in the summary data ($p = 0.13$ and $p = 0.17$, respectively).

The estimates of the differences on the intensive margin, reported in columns three and six of Table 1, reiterate what is seen in Figure 1(c). Without any controls in column 3 we see that although the Gift After treatment appears to gather larger positive donations than the other two conditions, the standard errors are relatively large. There is an appreciable change when we control for observables in column six, however. As one can see, there is a large effect of being on the warm list: previous donors give \$6.47 more than other participants ($p = 0.06$). Further, when we control for this effect, the point estimate on the Gift After treatment increases by \$0.37 and is now significant at the 10% level. One reason for this change is that there are fewer warm list participants in the Gift After treatment than in the other two conditions (11% compared to 14% and 15% in the other conditions). Although these differences in the number of previous donors are only on the verge of significance, the large donative effect of being on the

warm list makes a difference to the point estimates in Table 1.

All margins considered, we find that the sequencing of gift exchange matters in our field experiment. When the principal offers a gift after the agent, overall giving increases above that in the gift before treatment and a baseline without gifts. In addition, the improved performance of the gift after is partial due to participation - gifts after attract significantly more donations than the other two conditions - and partially due to the intensity of giving - gifts after increase the amount given by those who choose to donate.

4 Discussion

Given the symmetry of the basic gift exchange scenario modeled as a sequential prisoner's dilemma, it is not obvious that it should matter if the principal makes the first gift or the agent does. Perhaps because of this isomorphism, most theory and experiments model the principal making the first move simply because this is how the ordering was originally specified in the labor market context of Akerlof (1982). We do not assume that the sequencing of gift exchange is innocuous and test both orderings against a control in which no gifts are exchanged.

Our results are remarkable for a number of reasons. First, demonstrating that our results are not driven by our procedures, solicitation technique or location, we are able to replicate some important features of the nascent field experimental literature on gift exchange, especially the studies in which the principal moves first. Like both Falk (2007) and Alpizar et al. (2008), we find that when the principal offers the first gift donations increase but that the effect is modest and mostly on the extensive margin. Second, we also find that previously donating to a charity (i.e., being on the warm list) is a strong predictor of donating again. This replicates the results of Landry et al. (2010) and Carpenter and Matthews (2015) from the philanthropic field experimental literature.

What is most striking, however, is that our results seem to be more at odds with other recent studies in which the agent is asked to move first. As discussed above, in Newman and Shen (2012) and Chao (2016) the promise of conditional gifts, the process typically used by charities, seems to crowd out donations and yet this crowding does not jibe with our results (or the ubiquity of the practice among fundraisers). Our conditional, Gift After, treatment is actually more effective (on both margins) than

either the voluntary contribution mechanism or the unconditional gift.

As stated at the outset, the problem with comparing our results to the existing literature is that these studies do not consider both orderings of gift exchange. There is, however, one complementary study by Landry et al. (2012), also conducted in the charitable fundraising context and with a similar sample size, in which five treatments are studied: a standard voluntary contribution mechanism like ours and four other treatments that vary the size of the gift (small vs. large) and the conditionality of the gift (unconditional vs. conditional). This second aspect of the design is what makes the experiment close to ours because in some cases unconditional gifts are made to potential donors (i.e., the principal goes first) and in others gifts are only made to those who donate (i.e., the agent moves first). Landry et al. find that the effect of ordering depends, to some extent, on the size of the gift. Bearing in mind the setting that is most similar to ours, the one in which the gift is small, a token of appreciation so to speak, the principal-first ordering does a bit worse than the no gift control (about 9% worse) and the agent-first one does a bit better (4% better), a result that is somewhat similar to ours, at least concerning the ordering in which the agent goes first.

So where does this leave us? Overall, our results on the principal-first ordering look similar to what has been found in the literature while our agent-first results are more of an anomaly. One conjecture is that gifts after perform better because they identify donors to the public in a way not possible with the gift before sequencing. This is especially true if potential donors value social image (à la Benabou and Tirole, 2006). A second hypothesis to “square” the agent-first results and our lack of crowding has to do with the presentation of the gift. As stressed in Kube et al. (2012), taking the time to present an incentive as a gift matters and perhaps it matters in the sense that it makes the gift appear less instrumental or “controlling,” to use a phrase common in the intrinsic motivation literature. When gifts are small and more clearly designed as tokens of appreciation rather than some quid pro quo, it could be the case that the exchange is not soured and there is no crowding. That said, these are more suggestions for further experimentation than solid conclusions. In the end, our results with a small token of appreciation suggest that agent-first (or conditional) gift is optimal not only because it better incents reciprocity but because it also, by definition, is cheaper to implement - gifts are only exchanged when some surplus can be shared.

5 References

Abeler, J., Altmann, S., Kube, S., Wibral, M., 2010. Gift exchange and workers' fairness concerns: When equality is unfair. *Journal of the European Economic Association*, 8, 1299-1324.

Akerlof, G., 1982. Labor Contracts as Partial Gift Exchange. *Quarterly Journal of Economics*, 92, 543-569.

Akerlof, G., Yellen, J., 1986, *Efficiency wage models of the labor contract*. New York: Cambridge University Press.

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, 1047-1060.

Anderson, E., Simester, D., 2008. Does demand fall when customers perceive that prices are unfair? The case of premium pricing for large sizes. *Marketing Science*, 27, 492-500.

Andreoni, J., Payne, A., 2013, Charitable Giving. In: Auerbach, A., Chetty, R., Feldstein, M., Saez E. (Eds.). *Handbook of Public Economics*, Vol. 5. New York: Elsevier, 1-50.

Benabou, R., Tirole, J., 2006. Incentives and prosocial behavior. *American Economic Review*, 96, 1652-1678.

Carpenter, J., Matthews, P.H., 2015, Using Raffles to Fund Public Goods: Lessons from a field experiment. IZA Discussion Paper No. 8952.

Chao, M., 2016, Motivation crowding out in charitable giving: A field experiment. Williams College Working Paper.

Charness, G., Kuhn, P., 2011, Chapter 3 - Lab Labor: What Can Labor Economists Learn from the Lab? In: A. Orley, A., David, C., (Eds.). *Handbook of Labor Economics*. New York: Elsevier, 229-330.

Dufwenberg, M., Kirchsteiger, G., 2004. A Theory of Sequential Reciprocity. *Games and Economic Behavior*, 47, 268-298.

Falk, A., 2007. Gift exchange in the field. *Econometrica*, 75, 1501-1511.

Fehr, E., Klein, A., Schmidt, K., 2007. Fairness and contract design. *Econometrica*, 75, 121-154.

Gneezy, U., List, J., 2006. Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments. *Econometrica*, 74, 1365-1384.

Kahneman, D., Knetsch, J., Thaler, R., 1986. Fairness and the Assumptions of Economics. *Journal of Business*, 59, 285-300.

Karlan, D., 2005. Using Experimental Economics to Measure Social Capital and Predict Financial Decisions. *American Economic Review*, 95, 1688-1699.

Kleine, M., Kube, S., 2015, Communication and trust in principal-team relationships: Experimental evidence. Max Planck Institute for Research on Collective Goods Working Paper.

Kube, S., Marechal, M.A., Puppe, C., 2012. The Currency of Reciprocity: Gift Exchange in the Workplace. *American Economic Review*, 102, 1644-1662.

Landry, C., Lange, A., List, J., Price, M., Rupp, N., 2010. Is a donor in hand better than two in the bush? Evidence from a natural field experiment. *American Economic Review*, 100, 958-983.

Landry, C., Lange, A., List, J., Price, M., Rupp, N., 2012, Using donor gifts to drive fundraising: Theory and evidence from a natural field experiment. University of Georgia Working Paper.

Lee, S., Perrson, P., 2016. Financing from family and friends. *The Review of Financial Studies*, forthcoming.

Mellström, C., Johannesson, M., 2008. Crowding out in blood donation: was Titmuss right? *Journal of the European Economic Association*, 6: 845-63.

Newman, G., Shen, Y.J., 2012. The counterintuitive effects of thank-you gifts on charitable giving. *Journal of Economic Psychology*, 33, 973-983.