

You Owe Me[†]

By ULRIKE MALMENDIER AND KLAUS M. SCHMIDT*

In business and politics, gifts are often aimed at influencing the recipient at the expense of third parties. In an experimental study, which removes informational and incentive confounds, subjects strongly respond to small gifts even though they understand the gift giver's intention. Our findings question existing models of social preferences. They point to anthropological and sociological theories about gifts creating an obligation to reciprocate. We capture these effects in a simple extension of existing models. We show that common policy responses (disclosure, size limits) may be ineffective, consistent with our model. Financial incentives are effective but can backfire. (JEL C92, D62, D64, Z13)

Following Akerlof's (1982) seminal paper on gift exchange in labor markets, a large experimental literature has shown that gifts induce socially beneficial cooperation, both in the laboratory and in the field (Fehr, Goette, and Zehnder 2009). Much less attention has been devoted to the "dark side" of prosocial behavior: its negative externalities. Prosocial behavior toward one person may come at the expense of a third party. Consider the widespread use of business gifts. In a typical scenario, a procurement manager receives gifts (ranging from small tokens of appreciation, such as pens or coffee mugs, to bottles of wine or event tickets) from a supplier, who hopes to get favorable treatment. An extreme example is the pharmaceutical industry, estimated to spend \$8,000–\$15,000 per year on each physician in the United States, including luxurious dinners, conferences at attractive locations, and generous honoraria (Blumenthal 2004, p. 1885). Similarly, politicians and regulators receive gifts or campaign contributions from lobbyists. In all cases the recipient makes a decision on behalf of a third-party *client* who is often anonymous: shareholders, health insurances, or the electorate. Such practices have raised concerns, and stirred a regulatory debate, about the influence of gifts. Gift-giving has been blamed as a

*Malmendier: Department of Economics, University of California at Berkeley, 501 Evans Hall, Berkeley, CA 94720 (e-mail: ulrike@berkeley.edu); Schmidt: Department of Economics, University of Munich, Ludwigstrasse 28, D-80539 Munich, Germany (e-mail: klaus.schmidt@LMU.de). We would like to thank Yan Chen, Rachel Croson, Robert Dur, Ernst Fehr, Rudolf Kerschbamer, Matthias Kräkel, George Loewenstein, and Cristina Nistor. We would also like to thank seminar participants in the 2011 ASSA meetings, the Third Behavioral Economics Annual Meeting (BEAM), and at the universities of Aix-Marseilles, Berkeley, Emory, Innsbruck, Maastricht, MIT, Munich, Pennsylvania, Princeton, and Toulouse for comments and suggestions. David Arnold, Leonie Giessing, Jana Jarecki, Sara Kwasnick, Eddie Ning, Warren Ronsiek, and Slobodan Sudaric provided excellent research assistance. Financial support by Deutsche Forschungsgemeinschaft through SFB-TRR 15 and 190, the Excellence Initiative of the German government, and the Alfred P. Sloan Foundation is gratefully acknowledged.

[†]Go to <https://doi.org/10.1257/aer.20140890> to visit the article page for additional materials and author disclosure statement(s).

major contributor to weak corporate governance, to the dramatic rise of health care costs, and to wasteful pork barrel politics.¹ Nevertheless, these externalities have been largely ignored in the theoretical and experimental literature.

In this paper we use a controlled laboratory study to explore the negative externalities of gift-giving. First, we show that there is a powerful effect of the gift per se. Recipients reciprocate to gifts even if there are no monetary incentives for doing so, the gift is small, and it does not convey positive information about the gift giver or the product offered.² Recipients fully understand that the gift is given for self-ish reasons, namely to influence their behavior at the expense of a third party. We also observe a similar effect if a gift is expected, but not given: the decision maker discriminates against the potential gift giver. Second, we show that most standard models of social preferences have difficulties predicting the observed behavior. It is better predicted by anthropological and sociological theories positing that gifts create an obligation to reciprocate. We propose a simple extension of standard models of social preferences that formalizes this idea. Third, we conduct three policy treatments to evaluate the effectiveness of policy interventions that are often proposed to mitigate the effects of gifts with negative externalities.

In our main treatment (*Gift Treatment*), a decision maker has to buy one of two possible products on behalf of a client. Before making the decision, she may receive a small gift from one of the two producers. The gift is given unconditionally and before the producers learn about the value of their products, so that the gift cannot contain any information about the product's quality. The setting is anonymous, and players are rematched after each round. Neither client nor producers find out which product the decision maker picks, which excludes any reputation effects. Nevertheless, the effect of the gift is large. It is significantly stronger than in the classic gift-exchange game without a third party, which we conducted in a control treatment (*No Externality Treatment*). Furthermore, we observe that the gift is given with the intention to affect the decision of the recipient at the expense of a third party. We show that the recipients are fully aware of this intention, but they reciprocate nevertheless. We also find that, if a producer could have given a gift but chose not to do so, the decision maker often punishes him by refusing to buy his product even if that product is better. This punishment also comes at the expense of the client. Our results are robust to numerous variations, including making gift-giving more or less efficient, making the externalities on the client very large, and giving the client a more active role by *hiring* the decision maker.

Our experimental design allows us to test whether decision makers are aware of how strongly the gift affects their behavior. This question is much debated in practice. For example, a questionnaire study by Steinman, Shlipak, and McPhee

¹ See, e.g., Katz, Caplan, and Merz (2003); Blumenthal (2004); and Susman (2008). Policy initiatives range from voluntary codes of conduct (see, e.g., Murphy 1995 on corporate ethics statements and Grande 2010 on self-regulation in the pharmaceutical industry) to regulatory reforms and laws limiting the possibilities for gifts and requiring disclosure, such as the Lobbying Disclosure Act of 1995 and the Honest Leadership and Open Government Act of 2007 in the United States.

² In field settings, many gifts provide financial incentives or convey information. For example, physicians may prescribe more drugs of a pharmaceutical company after attending a conference sponsored by that company because they want to receive more sponsoring in the future or because of scientific information provided at the conference. These effects are important and explained by standard economic theory. In this paper, we focus on the effects of a gift per se, i.e., in the absence of any incentive or information effects.

(2001) found that only 39 percent of medical residents believe that gifts from pharmaceutical companies affect their own prescription behavior, but 84 percent believe that other physicians are influenced. In our experiment, we asked decision makers at the end to estimate how often their decisions coincided with the client's preferences, and used a quadratic scoring rule to induce unbiased estimates. On average, their estimates are highly accurate. They somewhat overestimate by how much others are affected by the gift, but the estimated difference between own and others' behavior is much smaller than in Steinman, Shlipak, and McPhee (2001).

The most prominent economic theories of other-regarding behavior have difficulties predicting the observed phenomena. Theories of outcome-based social preferences such as altruism (Andreoni and Miller 2002); maximin preferences (Charness and Rabin 2002); inequality aversion (Fehr and Schmidt 1999); and type-based reciprocity (Levine 1998) are refuted by the observed behavior. Theories of intention-based reciprocity (Rabin 1993, Dufwenberg and Kirchsteiger 2004) have multiple equilibria including one in which no gift is given and the decision maker always chooses the best product for the client. Thus, they have little predictive power. Furthermore, the underlying motives seem different. Producers admit in the questionnaire that they give the gift for selfish reasons, and the decision makers state that they perfectly understand this.

Our evidence suggests that a gift triggers an obligation to repay, independently of the intentions of the gift giver and the distributional consequences. It seems to create a bond between gift giver and recipient, in line with a large anthropological and sociological literature on gifts creating an obligation to reciprocate. In a classical study of small-scale societies, Mauss (1924) shows that reciprocity is the lubricant of social exchange and that people in these societies are under an obligation to give and to reciprocate. Prominent sociologists such as Gouldner (1960) and Blau (1964) argue that the obligation to reciprocate is a universal social norm.

We propose a simple theoretical framework that formalizes this idea. In the model the weight that player i attaches to the welfare of player j depends on the actions of j that affect i , relative to the expected behavior of j . A favorable act such as giving a gift strengthens the bond between the giver and the recipient, i.e., the weight of the giver's payoff in the recipient's utility function, and the recipient will reciprocate. The key difference to existing models of action-based reciprocity (e.g., Cox, Friedman, and Sadiraj 2008) is the prediction that the more the favorable act exceeds expectation, the stronger the positive response. We confirm this additional prediction in the data.

Finally, we conduct several policy treatments to evaluate remedies commonly proposed to mitigate the problem of gift-giving. In the disclosure treatment we inform clients which producer gave the gift and which product the decision maker chose. Decision makers' behavior remains very similar to the Gift Treatment without disclosure. This holds even if clients can punish decision makers (at a small cost to themselves) for choosing the wrong product.

Next, we address the effectiveness of size limits by varying the size of the gift. We find that larger gifts have *less* of an impact. When the gift is three times as large, decision makers choose the gift giver 22 percent *less* often than in the treatments where the gift is small. This is contrary to the logic of size limits, and may be surprising at first glance. But it is predicted by our model: decision makers correctly

expect that producers will send the gift with a higher probability if it benefits them more. The more the gift is expected, the smaller is the effect on behavior.

In a last policy treatment, the client can offer financial incentives to align the payoff of the decision maker with his own interest. This is highly effective. When the client shares profits, the effect of the gift is much less pronounced than in the Gift Treatment but still slightly stronger than in the No Externality Treatment. However, when clients do not offer profit sharing (but could have done so), the effect of the gift becomes even stronger than in the Gift Treatment. Again, these results are consistent with the model we propose.

The rest of the paper is organized as follows: the next section discusses related literature. Section II describes the experimental design. Section III shows that none of the most prominent economic theories of social preferences predict that the decision maker favors the gift giver. Section IV presents our main experimental results. We also analyze whether decision makers are aware of how gifts affect their behavior. Section V offers a theoretical framework to explain the observed behavior in terms of social preferences with endogenous reference groups. Section VI considers the policy treatments that test how to mitigate the effects of gift-giving. Section VII concludes.

I. Literature

In addition to the papers mentioned above (and the anthropological and sociological literature discussed in Section V), our paper is related to three branches of the economics literature. First, a large experimental literature, starting with Fehr, Kirchsteiger, and Riedl (1993), has established reciprocity as a motive for gift exchange. In almost all of the literature, the gift affects only the giver and the receiver; there are no externalities. A notable exception is the bribery game in Abbink, Irlenbusch, and Renner (2002) and Abbink (2004), where one player can bribe another player to take an action that is beneficial to him at the expense of all other players in the experiment. The authors show that repeated interaction sustains a bribery relationship, but the threat of a (probabilistic) penalty and staff rotation significantly weakens it. In contrast, we focus on the reciprocal effect of small gifts that are given unconditionally without any repeated interaction, penalties, or monetary incentives.

A second related strand of the literature studies gifts with positive or negative externalities in the field. Currie, Lin, and Meng (2013) conduct a field experiment in Chinese outpatient clinics, where two trained auditors acting as patients visit a physician in sequence. If the first patient gives a small gift to the physician, he receives better service and is less likely to be prescribed unnecessary and costly antibiotics than if no gift is given. Furthermore, if the first patient introduces the second patient as his friend, this patient also receives better service. Manacorda, Miguel, and Vigorito (2011) estimate the impact of a large antipoverty program in Uruguay on political support for the government. Those (quasi)randomly selected households who benefited from the program are 11–13 percentage points more likely to favor the government than those who did not benefit. These studies show that reciprocity in the presence of externalities is not restricted to the lab but extends to the field; however, they cannot identify the underlying behavioral mechanism.

Third, a large empirical literature studies business gifts, especially in the pharmaceutical industry. In a meta-study of 29 datasets, Wazana (2000, p. 373) concludes that gifts are “associated with increased prescription rates of the sponsor’s medication.” Campbell et al. (2007) conduct a survey of 3,167 physicians in 6 specialties and document the types of gifts given by the pharmaceutical industry and the nature of physician-industry interaction. Morgan et al. (2006) conducted a survey among physicians on whether it is ethical to accept gifts of the pharmaceutical industry and whether these gifts affect prescriptions. The general conclusion in this literature is that business gifts are widespread and effective. However, as discussed in Dana and Loewenstein (2003), the empirical literature cannot disentangle the causal factors that explain why gifts work.

II. Experimental Design

Our experimental design captures situations where one person has to rely on another person to make a decision on his behalf, and a third party has an interest in influencing this decision. We focus on the case where the latter party can give a small gift (such as a pen, coffee mug, or invitation for lunch), the gift is given unconditionally, and the parties interact only once. Such small gifts are common in many cultures and industries and, unlike bribes, are often legal and socially accepted.

In the experiment, two producers want to sell their products to a client. At the beginning of each period one producer is chosen randomly as the potential gift giver. This producer receives one additional token that he can pass on as a gift. The other producer cannot take any action. The potential gift giver, called producer *X*, offers product *X*; the other producer, called producer *Y*, offers product *Y*. The client has to buy either product *X* or product *Y* but has to rely on an expert to make this decision on his behalf. We call the expert the decision maker (DM). DM receives a fixed wage of 20 tokens for her services and is instructed to choose the best product for the client (“take a decision that is in the best interest of the client”). If she chooses product *X* (*Y*, respectively), producer *X* (*Y*) receives a positive (quasi-)rent of 16 tokens, while the other producer gets 0. Before DM decides, producer *X* can pass the additional token on to her, in which case it doubles and DM receives 2 tokens. The gift is unconditional, cannot be refused, and there is no repeated interaction—subjects are anonymously rematched after each round. The client is aware of the possibility of a gift, but does not know whether a gift was actually given, nor which product was chosen.

A typical session has 24 subjects: 6 DMs, 6 clients, 6 producers *X*, and 6 producers *Y*.³ There are 20 periods. In each period DM is anonymously matched with a new client and new producers *X* and *Y*. Products *X* and *Y* are simple 50/50 lotteries. The payoffs are natural numbers between 3 and 20. For example, product *X* may pay 5 or 11 with equal probability, while product *Y* pays 3 or 17. The lottery pairs in each period are fixed throughout all treatments and sessions. In four periods, the expected value (EV) of lottery *X* exceeds the expected value of *Y* by 2; in six periods, the two lotteries have the same expected value (but differ in variance); in six

³There is one session in the Incentive Treatment in which only 20 subjects participated (5 DMs, 5 clients, 10 producers). In the No Externality Treatments there are no clients, so in these sessions we had 8 DMs and 16 producers.

periods, the expected value of X is 2 points lower than the expected value of Y ; and in four periods, the expected value of X is 6 points lower. In this last case, product Y first-order stochastically dominates product X , i.e., every rational decision maker (no matter how risk-averse or risk-loving) prefers Y to X . (There are also three cases with an expected-payout difference of 2 where Y first-order stochastically dominates. Online Appendix-Table A1 in online Appendix B shows all 20 lottery pairs.)

Note that Producer X has to decide whether to send the gift before he learns what the products X and Y are in this period. Thus, the gift cannot signal product quality. The producers never learn which product DM chooses. They are only informed about their total payoff after all 20 periods. Thus, there is no learning about the effectiveness of gifts, and a producer's future behavior cannot be affected by choosing or not choosing his product.

DM learns who the potential gift giver is and whether he sent the gift. She then sees the two lotteries and chooses one for her client. Her payoff is unaffected by her choice, and she does not learn how the lotteries resolve. The client does *not* know who the potential gift giver is and whether he sends the gift. He sees the two products and is asked which one he would choose if, hypothetically, he could decide himself. He does not observe which product DM chooses, nor the outcome of the lottery. At the end of the experiment he only learns the sum of his payoffs in all 20 periods.

In each session the instructions are read aloud. After 20 periods, subjects are asked to answer a questionnaire. In the first part, DMs estimate how often their own decision and the decision of the other DMs coincided with the preferred product of the clients. Similarly, clients and producers estimate how often the decision makers chose the product that the clients would have preferred. The answers are incentivized with a quadratic scoring rule. In the second part, we ask subjects about the motives for their own decisions and their beliefs about the motives of the other subjects.

We compare the results of this *Gift Treatment* (GT) to two other treatments: the *Baseline Treatment* (BT) and the *No Externality Treatment* (NET). In both of these treatments, we use exactly the same lotteries (*products*) in the same sequence as in the Gift Treatment.

In the *Baseline Treatment*, producers cannot send a gift and gifts are never mentioned. This treatment shows whether decision makers choose the products preferred by their clients if nobody tries to influence them. Comparing BT and GT allows us to test both for the effect of gift-giving and for the effect of not giving a gift (despite having the option) relative to a world without gifts.

In the *No Externality Treatment*, there is no client and no fixed wage for the decision maker. DM buys the product for herself and is full residual claimant of the lottery payoffs. This treatment allows us to estimate to what extent the effect of gifts in the Gift Treatment reflects the fact that DM acts on behalf of a third party and does not bear the consequences of her decisions.

We test the robustness of our results in several control treatments (discussed in Section IVD). Furthermore, we investigate three policy treatments: disclosure with and without punishment, variation in gift size, and financial incentives (described in Section VI).

Discussion of Design Features.—The experimental design has several features that merit discussion. First, only one of the two producers can make a gift. Had we

allowed both producers to offer gifts, both producers would likely have done so in most periods, in which case we would have lost many observations. Alternatively, we could allow both producers to also choose the size of their gifts and to enter a rent-seeking tournament which may result in higher gifts. Or, we could vary the externalities of gifts for the other producer. These are interesting avenues for future research, and the restriction to a noncompetitive setting has to be kept in mind for the interpretation of our results.

Second, in the experiment the decision maker had no choice but to accept the gift. We chose this experimental design to isolate the pure effect of the gift without further complications that arise if the gift giver does not know whether his gift is accepted. In reality it is often difficult to refuse the gift, but not impossible. Whether and when decision makers refuse to accept gifts that are intended to influence their behavior are interesting questions for future research.

Third, the products are simple lotteries over monetary outcomes. This design feature gives the decision maker some moral wiggle room: even when product *Y* has a higher expected value than product *X*, DM may justify buying *X* because it is less risky or because it gives a higher payoff in one state of the world. Note, though, that our design includes cases in which every rational agent prefers *Y* because it first-order stochastically dominates *X*. The Baseline Treatment will show that subjects had no difficulties evaluating the simple lotteries. Also note that many real-world decisions that motivate our study feature probabilistic consequences, such as the effects of a drug on a patient or a policy measure on the general public.

Fourth, the client and the producers do not observe the decision of the decision maker but learn only the sum of their payoffs over all 20 periods at the end of the experiment. We do not impose this design feature for realism. Rather, we use it to make sure that reputation building or learning cannot affect our results, but that it is the effect of the gift per se that we observe.

Experimental Procedures.—We conducted 31 sessions with 20–24 participants in each session at the MELESSA laboratory of the University of Munich in 2010, 2011, and 2015. Subjects were undergraduate students of various disciplines from the University of Munich and the Technical University of Munich.⁴ We involved 740 subjects, generating a dataset of 4,140 observations.⁵ The vast majority (93 percent) of students were in the typical age range of 20–29, and slightly more than one-half (54 percent) were women. Upon arrival at the lab, subjects were randomly and anonymously assigned to the different roles. Sessions lasted about one hour. On average, subjects earned €14 (\$19.15), which includes a show-up fee of €4 (\$5.47). Further summary statistics are in online Appendix Table A2.

⁴We used the software ORSEE (Greiner 2004) for recruitment and z-Tree (Fischbacher 2007) for the experiments.

⁵One additional session containing only 20 subjects (Disclosure Treatment) had to be excluded ex post because of a problem in the matching procedure in z-Tree. All results are robust to including this session.

III. Behavioral Predictions

To guide our empirical analysis, we derive the predictions of existing theories of social preferences for the response to (not) receiving a gift. Remember that the gift is given unconditionally and prior to her decision. Thus, the traditional model of rational self-interested behavior does not predict that DM favors the gift giver. She is indifferent between the two products, no matter whether the gift was given. If there is no client and DM chooses the product for herself, the traditional model predicts that she chooses the better product. Thus, if we want to explain why the gift systematically affects her decision, we have to look for alternative models.

First, we consider outcome-based social preferences, and in particular the three forms that have received most interest in the literature: altruism, maximin preferences, and inequality aversion. Suppose that DM has outcome-based social preferences $U^{DM}(m^{DM}, m^X, m^Y, m^C)$, where m^i is the expected monetary payoff of player $i \in \{DM, X, Y, C\}$ and U^{DM} is invariant to permutations of (m^X, m^Y, m^C) . Under (i) altruism in the form of utilitarianism (Andreoni and Miller 2002), DM's utility increases with the sum of the payoffs of the others; under (ii) maximin preferences (Charness and Rabin 2002), DM's utility increases with the payoff of the worst-off in the group; and under (iii) inequality aversion (Fehr and Schmidt 1999; Bolton and Ockenfels 2000), DM dislikes to be worse off and (to a lesser degree) to be better off than the other players.

We assume that the decision maker is risk neutral and evaluates products X and Y by their expected values.⁶ We say that DM *favors* a producer if she chooses his product no matter how it compares to the other product. We say that DM favors the client if she chooses the product with the higher expected value or, if both have the same expected value, with the smaller variance. We assume that if DM is indifferent she favors the client, as frequently imposed in contracting games with multiple equilibria. This is confirmed by the results of the Baseline Treatment.

All outcome-based theories predict unambiguously that DMs should not be influenced by gifts but should maximize the expected utility of their clients.

PROPOSITION 1: *Suppose that the decision maker is motivated by (i) altruism (utilitarianism), (ii) maximin preferences, or (iii) inequality aversion. Then, we have:*

- (a) *In the Baseline Treatment, where no gift can be passed on, DM always favors the client.*
- (b) *In the Gift Treatment, if producer X did pass on the gift, DM always favors the client.*
- (c) *In the Gift Treatment, if producer X did not pass on the gift, DM favors the client if she is altruistic (utilitarian) or inequality averse, but favors producer Y if she has maximin preferences.*

⁶Most existing social-preference theories do not explicitly consider choices between lotteries. Since the experimental stakes are fairly small, risk aversion should not affect decision making and, to a first approximation, risk neutrality is not restrictive. Note also that in our experiment the decision maker never observes the outcome of the lotteries.

PROOF:

See online Appendix A.

In other words, Proposition 1 implies that DM never favors X , no matter whether a gift is given or not. This is contradicted by the experimental results we show below.

It is instructive to briefly go over the main arguments of the proof. DM cannot affect the payoff distribution of producers: in the Baseline Treatment and in the Gift Treatment, if X has sent the gift, one producer gets 16, and one gets 0, regardless of the product chosen. But she can affect the expected payoff of the client. Thus, and given that DM’s payoff is (weakly) higher than that of any other player in any state of the world, all three outcome-based theories of social preferences predict that DM favors the client. If producer X has not sent the gift, then the payoff distribution of producers X and Y is $(17, 0)$ if DM chooses X and $(1, 16)$ if she chooses Y . For an altruistic (utilitarian) or inequality averse DM this does not matter; she still maximizes the client’s expected payoff. Under maximin preferences, she maximizes the payoff of the worst off and favors producer Y .

Second, we consider type-dependent preferences, as in Levine (1998). Assume for simplicity that there are two types of players, *kind* and *selfish* types. Kind types care positively about the payoffs of other players who are also kind, but they do not care about payoffs of selfish players. Selfish types care only about their own payoffs. A player’s type is private information. It is common knowledge that the ex ante probability of being kind is μ , i.e., $\mu^{DM} = \mu^X = \mu^Y = \mu^C = \mu$, with $0 < \mu < 1$ for all $i \in \{DM, X, Y, C\}$. Let μ_i^j denote the (updated) belief of player i about the type of player j , with $i, j \in \{DM, X, Y, C\}$ and $i \neq j$. Then the expected utility of a kind player i is

$$(1) \quad U^{i,kind} = m^i + \sum_{j \neq i} \mu_i^j \cdot \alpha \cdot m^j,$$

where $\alpha > 0$ is the (common) degree to which a kind player i cares about the payoff of a kind player j .⁷ The utility function of a selfish player i simply is $U^{i,selfish} = m^i$.

PROPOSITION 2: *Suppose that the decision maker has type-dependent preferences. In any pooling equilibrium, the kind and the selfish type of DM favor the client. Any (partially) separating equilibrium requires that the probability of the decision maker choosing product X does not increase by more than 1/16 when the gift is given compared to when the gift is not given.*

PROOF:

See online Appendix A.

Proposition 2 implies that models of type-based reciprocity are not consistent with our data either. In our experiment, DMs favor producer X if the gift was given and

⁷The results remain qualitatively the same if we assume that a kind player cares about the payoff of a selfish player to the degree $\underline{\alpha}$ and about the payoff of a kind player to the degree $\bar{\alpha}$, with $\underline{\alpha} < \bar{\alpha}$ and $(1 - \mu)\underline{\alpha} + \mu\bar{\alpha} > 0$.

producer Y if the gift was not given. In any pooling equilibrium, instead, DMs should favor the client. And in any (partially) separating equilibrium, the gift is a signal of kindness only if the reaction of DMs to the gift is not too large (less than $1/16$). In the data shown below, the probability that DM chooses X increases by 44 percentage points if the gift is given, much more than $1/16 = 6.25$ percent. Thus, giving the gift does not signal a kind type as a selfish producer has a strong incentive to mimic the kind type.

Finally, we consider intention-based reciprocity (Rabin 1993; Ruffle 1999; Dufwenberg and Kirchsteiger 2004). To apply it to our experiment, we simplify DM's strategy space to choose only between action X , i.e., choosing product X , and action C , i.e., favoring the client. We use the notion of sequential reciprocity equilibrium (SRE) of Dufwenberg and Kirchsteiger (2004).

PROPOSITION 3: *Suppose that DM and producer X exhibit intention-based reciprocity. If they care strongly enough about intentions, then there exists an SRE in which producer X sends the gift and DM chooses X ; but there also exists an SRE in which X does not give the gift and DM chooses C .*

PROOF:

See online Appendix A.

Psychological games with intention-based reciprocity are consistent with many interesting phenomena, but are also plagued by multiple equilibria. It is an equilibrium that both players are kind to each other if both expect the other player to be kind as well. Here, all producers send the gift and DMs favor the gift giver. But it is also an equilibrium that both players behave unkindly if they believe that the other player is hostile. Here, no gifts are given and DMs favor their clients. Thus, intention-based reciprocity is consistent with the main findings, but lacks predictive power.⁸

IV. The Effect of Gifts

A. Baseline without Gift-Giving

Decision makers are instructed to choose in the best interest of their clients. Before we study how gifts of third parties affect this decision, we establish what happens when there are no gifts. Which products do decision makers choose, and which products do clients prefer?

In Figure 1, we show how often product X is preferred by the client (dark bars) and how often it is chosen by DM (light bars), both on average over all 20 periods and separately for cases when the expected value of X is higher than, equal to, or lower than that of Y . These choices serve as the benchmark for all other treatments, where producer X has the option of giving a gift.

The figure shows that clients strongly prefer the lottery with the higher expected value: 97 percent prefer product X if it has the higher expected value. If the expected

⁸ A related approach is "guilt aversion" (Charness and Dufwenberg 2006): people feel guilt if they do not meet others' expectations. This approach is not applicable since neither producer nor client learn which product DM chooses.




FIGURE 1. RELATIVE FREQUENCIES OF CHOOSING X IN THE BASELINE TREATMENT

value of lottery X is, instead, 2 points smaller than that of Y, only 8 percent of clients prefer X. And if X has a disadvantage of 6 points (and is first-order stochastically dominated by Y), no client prefers X.

The choices of decision makers are closely aligned with the preferences of clients. The overwhelming majority chooses the lottery with the highest expected value. There is no statistically significant difference between their choices and the preferred choices of the clients, both when expected values differ and when they are equal.⁹ Note that clients prefer and DMs choose X in 35–36 percent of all cases, in line with product X having the higher expected value (or the same expected value and the lower variance) in 7 out of 20 periods, i.e., in 35 percent of all cases.

Also note that a little less than 50 percent of clients and DMs choose X when expected values are equal. Out of the six periods in which expected values are the same, X has a larger variance in three periods and a smaller variance in the other three. About two-thirds of subjects (72.2 percent of clients and 63.2 percent of DMs) choose the product with the smaller variance. Hence, in addition to a strong preference for higher expected values, we observe a weaker preference for the lower variance. We ignore the (insignificant) effects of the variance in the following.

B. Gift-Giving with Externalities

What happens if producer X can make a gift? Figure 2 compares the choices of DM in the Baseline Treatment (middle bars) to the Gift Treatment when the gift was sent (bars to the left) and when it was not sent (bars to the right). Note that the potential gift givers sent the gift in 73.1 percent of all cases (702 out of 960), and 46.9 percent *always* sent the gift, while 7.3 percent *never* did so.¹⁰

The first set of bars shows that the average effect of the gift is very large: if the gift is given, X is chosen almost twice as often as in the BT, with 65.0 rather than 36.0 percent frequency. If the gift is not given, the frequency decreases by

⁹A Wilcoxon-Mann-Whitney test comparing the choices of DMs and the preferred choices of clients does not reject the hypothesis that the two are drawn from the same distribution, both when the expected values of X and Y differ ($p = 0.876$) and when they are equal ($p = 0.410$).

¹⁰Aggregating over all treatments, in which gift-giving is possible we observe that 40.5 percent of producers play a pure strategy, most of which (34.1 percent of the total number of producers) always give the gift.




FIGURE 2. RELATIVE FREQUENCIES OF CHOOSING X IN THE GIFT TREATMENT
(With Comparison to the Baseline Treatment)

more than one-third, relative to the BT, to 20.9 percent. Fligner-Policello robust rank-order tests comparing the average decisions of each DM in the GT when the gift was given (not given) and the BT show that these differences are highly significant (both with $p < 0.001$).¹¹ In other words, DMs strongly reciprocate to the gift by favoring the gift giver and exert negative reciprocity if a potential gift giver chose not to send a gift.

The next set of bars shows that, if X has a two point advantage in terms of expected value, it is almost always chosen, both in the Baseline and in the Gift Treatment when the gift is given (88.5 and 93.7 percent, respectively). However, when producer X does not pass on the gift, the fraction of DMs buying X drops to 54.0 percent even though it is the better product. 46.0 percent of DMs *punish* producers for not giving the gift, at the expense of the client.

If the expected value of the two products is the same, we observe both the positive and the negative effect of the gift: the fraction choosing X roughly doubles, compared to the Baseline, if the gift was sent, from 43.8 to 82.6 percent, but almost halves, to 23.5 percent, if the gift was not sent. Finally, if X has a lower expected value, it is rarely chosen in the Baseline Treatment (10.2 percent) and even less often in the Gift Treatment when the gift was not sent (8.1 and 2.4 percent for 2 and 6 points disadvantage, respectively). However, if producer X did send the gift, his product is chosen in 47.0 percent of cases when the expected value is 2 points lower, and in 37.8 percent of cases when it is 6 points lower.¹² In fact, when considering only cases where X is first-order stochastically dominated (which includes all four cases with a 6 point disadvantage, and three cases where X has a 2 point disadvantage), we

¹¹ We use the Fligner-Policello robust rank order-test rather than the Wilcoxon-Mann-Whitney test to account for unequal variances across populations. This is important in our analysis as the number of periods we average over varies. For example, in the Baseline Treatment, we calculate the frequency of choosing X over 20 periods. But in the Gift Treatment we calculate separately average decisions when the gift was given and when it was not given and hence use fewer periods; the variance is mechanically larger.

¹² Fligner-Policello robust rank-order tests comparing the average decisions of DMs in the GT when the gift was given to their average decisions in the BT show significant differences if X has the same expected value ($p < 0.001$), a 2 point disadvantage ($p < 0.001$), or a 6 point disadvantage ($p = 0.018$). When the gift was not given, the difference is statistically significant if X has a 2 point advantage ($p < 0.005$) or the same expected value ($p < 0.001$).

TABLE 1—GIFT TREATMENT

	OLS (1)	OLS (2)	Logit (3)	OLS (4)	OLS (5)	Logit (6)	OLS (7)	OLS (8)
Gift given	0.289 (0.028)	0.290 (0.028)	0.290 (0.028)				0.274 (0.028)	
Gift not given	-0.151 (0.034)	-0.152 (0.033)	-0.183 (0.040)				-0.138 (0.033)	
Product <i>X</i> has higher EV [+2]				0.448 (0.053)	0.452 (0.053)	0.495 (0.053)		0.452 (0.053)
Product <i>X</i> has lower EV [-2]				-0.333 (0.049)	-0.337 (0.049)	-0.425 (0.057)		-0.337 (0.049)
Product <i>X</i> has lower EV [-6]				-0.333 (0.052)	-0.333 (0.052)	-0.401 (0.058)		-0.333 (0.052)
(Product <i>X</i> has same EV) × (gift given)				0.389 (0.049)	0.389 (0.049)	0.404 (0.042)		0.391 (0.049)
(Product <i>X</i> has same EV) × (gift not given)				-0.203 (0.063)	-0.204 (0.062)	-0.216 (0.061)		-0.211 (0.062)
(Product <i>X</i> has higher EV [+2]) × (gift given)				0.051 (0.039)	0.053 (0.039)	0.163 (0.113)		0.043 (0.039)
(Product <i>X</i> has higher EV [+2]) × (gift not given)				-0.345 (0.078)	-0.349 (0.078)	-0.361 (0.050)		-0.335 (0.075)
(Product <i>X</i> has lower EV [-2]) × (gift given)				0.366 (0.044)	0.367 (0.044)	0.439 (0.050)		0.373 (0.043)
(Product <i>X</i> has lower EV [-2]) × (gift not given)				-0.023 (0.039)	-0.024 (0.039)	-0.069 (0.115)		-0.025 (0.039)
(Product <i>X</i> has lower EV [-6]) × (gift given)				0.273 (0.051)	0.272 (0.051)	0.369 (0.066)		0.272 (0.051)
(Product <i>X</i> has lower EV [-6]) × (gift not given)				-0.080 (0.040)	-0.075 (0.040)	-0.309 (0.150)		-0.075 (0.040)
(Product <i>X</i> has lower EV[-16]) × (gift given)								-0.361 (0.065)
(Product <i>X</i> has lower EV[-16]) × (gift not given)								-0.255 (0.110)
Female		0.036 (0.026)	0.042 (0.030)		0.038 (0.022)	0.060 (0.035)	0.044 (0.025)	0.044 (0.021)
Economist		0.021 (0.026)	0.024 (0.030)		0.021 (0.022)	0.034 (0.035)	0.025 (0.026)	0.021 (0.022)
Period		-0.004 (0.002)	-0.005 (0.002)		0.002 (0.002)	0.004 (0.003)	-0.006 (0.002)	0.002 (0.002)
Constant	0.360 (0.022)	0.379 (0.037)		0.438 (0.042)	0.387 (0.048)		0.400 (0.034)	0.386 (0.047)
Observations	1,440	1,440	1,440	1,440	1,440	1,440	1,560	1,560
(Pseudo) R^2	0.128	0.131		0.371	0.373		0.119	0.382
Sample		GT; BT	GT; BT		GT; BT		GT; LET; BT	

Notes: The GT, BT, and LET samples contain all data from the Gift, Baseline, and Large Externality Treatment, respectively. *Gift given* indicates that producer *X* sent the gift in the Gift Treatment. *Gift not given* indicates that producer *X* did not send the gift in the Gift Treatment. Robust standard errors are reported. The interaction term (Product *X* has lower EV[-6]) × (gift not given) perfectly predicts the outcome in logit regression (5).

find that *X* is chosen 38.4 percent of the time when the gift was given—even though there is no ambiguity about the client’s preferred product.

In Table 1, we evaluate the statistical significance of the effects in a controlled regression framework. We pool the BT and GT data and regress an indicator variable for the choice of *X* on dummies for *Gift given* and *Gift not given* (in the Gift

Treatment) using both a linear probability model (ordinary least squares (OLS)) and maximum-likelihood estimation (logit). The regression framework allows us to control for demographics and period effects and to account for heteroskedasticity, within-subject, and within-session correlation. We use robust, cluster-robust (clustering by subject), and wild cluster bootstrap- t procedures (clustering by session). In Table 1, we report the regression with robust standard errors. Online Appendix Table A3 displays the bootstrap results for session-level clustering.¹³

Columns 1 to 3 of Table 1 show the average effect of the gift, corresponding to the first set of bars in Figure 2. Regardless of the econometric model, giving a gift almost doubles the probability of the gift giver's product being chosen, and not giving a gift (when gift-giving is possible), reduces the probability by more than 40 percent. Both effects are highly significant ($p < 0.001$).¹⁴ Adding controls does not affect the coefficient estimates. In particular, neither gender nor the field of study (economics/business versus others) has a significant effect.

In the next three columns, we refine the regression model to account for the differences in expected values of the two products, as we did in the second to fifth set of bars in Figure 2. The intercept captures the case of "no difference in expected value" in the BT. The size and standard error of its coefficient estimate, as well as those of all indicators for expected-value differences in the BT ("Product X has higher/lower EV") go in the expected direction: better products are almost always chosen (with a probability of more than 90 percent), and worse products almost never (with a probability of less than 10 percent). All effects are highly significant. Hence, in a world without gifts, the influence of product value is economically and statistically highly significant.

The interaction terms in the rows below confirm the strong influence of gifts illustrated in Figure 2. All positive and negative effects discussed above are statistically significant, regardless of the econometric model. The insignificant or less significant coefficient estimates for $(EV [+2]) \times (\textit{Gift given})$ and $(EV [-2]) \times (\textit{Gift not given})$, and $(EV [-6]) \times (\textit{Gift not given})$ are cases of left- or right-censoring: product X is chosen almost always or almost never already in the Baseline Treatment. Thus, passing or not passing the gift cannot significantly alter the probabilities.

Given the powerful effect of first-order stochastic dominance (FOSD) in the baseline treatment, we also reestimate our model including control variables for first-order stochastic dominance. We find that the FOSD effect is large and significant when we do not control for differences in expected payoff. For example, we estimate a coefficient of 0.257 (standard error 0.0187) and of 0.346 (standard error 0.0230) in the model specifications of columns 1 and 2 of Table 1. However, the coefficient becomes insignificant and shrinks by 80–90 percent when we control for differences in expected payoffs of X and Y , as done in the specifications in

¹³We cannot cluster by session on account for within-session correlations because the number of clusters is too small. We implement the wild cluster bootstrap- t procedure of Cameron, Gelbach, and Miller (2008) to deal with cases of few clusters. We compute the significance levels using 10,000 repetitions. For completeness (and also because the number of DMs is sometimes too small to apply standard asymptotic tests), we also wild cluster bootstrap assuming DM clusters. The latter results are very similar to cluster-robust standard errors. We also reestimated the coefficients using fixed effects, both with robust standard errors and with clustered standard errors.

¹⁴Cluster-robust standard errors and wild cluster bootstrapping for individual-level clusters yield similar results. When bootstrapping with session-level clusters, the effects of (not) giving are significant (both with $p < 0.001$) in the baseline specification (online Appendix Table A3, column 1) and when adding controls (column 2).

columns 4 to 9 of Table 1. This finding suggests that there is no significant effect of the moral wiggle room embedded in nonstochastically dominated choices above and beyond the expected-payoff differences.

Overall, the Gift Treatment shows that gifts can have strong externalities. Clients receive the worse product with 43 percent probability (instead of 10 percent) if the producer of the worse product has sent a gift or the producer of the better product chose not to send a gift.

C. Gift-Giving without Externalities

We now ask how the effects of gifts compare to a setting without externalities. If gifts have the same effect, regardless of who pays the cost, there is no obvious distortion and it is more difficult to argue that gifts induce inefficient behavior. If, instead, DMs act differently on their own account, the possibility of gift-giving is distortive and likely to be welfare reducing for third parties.

There is no client in the *No Externality Treatment* (NET). DM decides on her own behalf and, instead of receiving a fixed wage, is full residual claimant. Figure 3 compares the BT (middle bars) to the NET when the gift was sent (left bars) and when it was not sent (right bars). Note that 49.1 percent of the potential gift givers (157 out of 320 cases) decide to send the gift.

Figure 3 shows that, even when acting on their own behalf, DMs choose the potential gift giver more often when he gives the gift (54.8 percent) than when he does not (28.8 percent) or when there is no possibility of giving a gift (36.0 percent in the BT). However, the effects are weaker than in the GT, where we observed an increase to 64.9 percent upon receiving the gift and a reduction to 20.9 percent upon not receiving a possible gift. The columns to the right confirm that DMs punish less often for *not* giving than in the GT in Figure 2. Also, when product *X* is much worse (6 point disadvantage in expected value) the influence of the gift vanishes.¹⁵

These visual impressions are confirmed by the estimations reported in Table 2. For brevity, we only show the linear probability model; the logit results are very similar in terms of economic and statistical significance. All specifications include the controls for gender, economics major, and period. The average effects are reported in columns 1 to 3. We use the pooled data of BT and NET in column 1, and we add the GT data in columns 2 and 3. We observe a significant but smaller increase in the frequency of choosing *X* when the gift is sent. Both the increase itself (amounting to 18.4–18.7 pp) and the difference in increase relative to the GT (–10.3 pp) are highly significant. The decrease after not receiving the gift (–6.8 to –7.1 pp) is marginally significant in the BT/NET sample (column 1) and insignificant in the BT/GT/NET sample (column 2), while the 8.4 pp difference in decrease relative to GT (column 3) is marginally significant.¹⁶ When we account for expected-value differences, we see that the effect of the gift is significantly positive in all cases but

¹⁵Fligner-Policello robust rank-order tests comparing the reveal that we cannot reject the hypothesis that the DMs' decisions in GT and NET when the gift was sent are drawn from the same distribution ($p = 0.198$). When the gift was not sent, they are significantly different ($p = 0.032$). Subsampling by EV differences, decisions are different when the gift was given and *X* has a 6 point disadvantage ($p = 0.023$) and when the gift was not given and *X* has a 2 point advantage ($p = 0.079$).

¹⁶Table A4 in the online Appendix shows that session-level clustering yields very similar results.




FIGURE 3. RELATIVE FREQUENCIES OF CHOOSING X IN THE NO EXTERNALITY TREATMENT (With Comparison to the Baseline Treatment)

TABLE 2—NO EXTERNALITY TREATMENT

	Diff. to BT (1)	Diff. to BT (2)	Diff. to GT (3)	Diff. to BT (4)	Diff. to BT (5)	Diff. to GT (6)
NET: gift given	0.184 (0.046)	0.187 (0.046)	-0.103 (0.044)			
NET: gift not given	-0.071 (0.043)	-0.068 (0.042)	0.084 (0.044)			
(Product X has same EV) × (NET: gift given)				0.322 (0.071)	0.325 (0.071)	-0.064 (0.064)
(Product X has same EV) × (NET: gift not given)				-0.084 (0.085)	-0.080 (0.084)	0.124 (0.087)
(Product X has higher EV [+2]) × (NET: gift given)				0.112 (0.034)	0.115 (0.034)	0.062 (0.022)
(Product X has higher EV [+2]) × (NET: gift not given)				-0.061 (0.074)	-0.058 (0.074)	0.290 (0.097)
(Product X has lower EV [-2]) × (NET: gift given)				0.195 (0.074)	0.195 (0.073)	-0.171 (0.076)
(Product X has lower EV [-2]) × (NET: gift not given)				-0.026 (0.047)	-0.020 (0.047)	0.003 (0.049)
(Product X has lower EV [-6]) × (NET: gift given)				-0.066 (0.050)	-0.063 (0.050)	-0.335 (0.055)
(Product X has lower EV [-6]) × (NET: gift not given)				-0.106 (0.033)	-0.103 (0.032)	-0.027 (0.025)
Dummies for (GT: gift given) and (GT: gift not given)		X				
Dummies for (gift given) and (gift not given)			X			
Dummies for EV differences				X	X	X
Dummies for EV differences interacted with (GT: gg) and (GT: gng)					X	
Dummies for EV differences interacted with (gg) and (gng)						X
Controls for gender, major, and period	X	X	X	X	X	X
Sample	NET, BT	NET, GT, BT	NET, BT	NET, BT	NET, GT, BT	NET, GT, BT
Observations	800	1,760	800	800	1,760	1,760
R ²	0.037	0.123	0.123	0.437	0.397	0.397

Notes: The NET sample contains all observations from the No Externality Treatment; the BT sample all observations from the Baseline Treatment; and the GT sample all observations from the Gift Treatment. The abbreviation gg indicates “gift available and given” and gng indicates “gift available but not given.” Constant included. Robust standard errors are reported.

the expected-value difference of -6 (column 4). The same is true when we add the observations from the GT and control for the corresponding interactions “GT: Gift given” and “GT: Gift not given” (column 5). Relative to the GT, column 6 shows that there is significantly less punishment for not sending the gift if X is the better product and significantly less reward for sending the gift if X is much worse.

We draw two main conclusions. First, a gift is effective even when DM acts on her own behalf, consistent with large experimental evidence on gift exchange without externalities and with many firms offering their customers small gifts. If price or quality differences to competing products are not too large, gifts tip the balance and induce customers to buy from the gift giver. Second, if there are no externalities, the effect of the gift disappears when the product is much worse than the competing product. For example, when the gift is given, 37.7 percent of all decision makers in the GT choose X even though it has a 6 point disadvantage after receiving the gift, but less than 5 percent do so in the NET. Thus, the differences in behavior are particularly strong when the external effect is large. This suggests that gift-giving can have a large negative impact on social welfare.

D. Robustness Checks

Efficiency of the Gift.—In the GT, the gift increases social surplus: it costs the gift giver 1 point while the benefit to DM is 2 points. To test whether this design feature affects the behavior of producers and DMs, we conduct a control treatment (GT2:2) in which the cost is 2, so the gift is just a transfer. We find that, in this treatment, the gift is given in 67.9 percent of all cases, which is similar to 73.1 percent in GT. Figure 4 suggests that DM’s behavior remains quite similar to that in GT. The regression analysis in columns 1 and 2 of Table 3 confirms that the differences to BT are similar to those of the GT, and there are no significant differences between GT and GT2:2.

We also conduct a control treatment (GT2:1) in which gifts are inefficient: the cost of the gift is 2 points and the benefit only 1 point. Hence, relative to its benefits, the cost to the producer is four times as large as in the original GT. Nevertheless, the gift is sent in 63.8 percent of cases, only about 10 pp less than in GT. Panel A of Figure 4 shows that if the gift is given, DMs choose X slightly less often (58.2 percent) than in GT (65.0 percent). Moreover, if the gift is not given, X is chosen more often than in GT (29.9 percent compared to 20.9 in GT). It seems that DMs do not expect producers to give the gift and do not punish them as much if gift-giving is costly and inefficient.

Overall, the additional treatments confirm the robustness of our results to variations in the efficiency of giving. They also point to other determinants of the response to gifts. In particular, we note that DMs react slightly less strongly to receiving (or not receiving) the gift in GT2:1 than in GT. Possible explanations include the possibility that DMs dislike inefficiency or care about the absolute benefit of gift-giving to them.¹⁷

¹⁷The latter explanation is in line with the finding in Pan and Xiao (2016), who aim to disentangle the roles of costs and benefits of a gift in triggering reciprocity. In their setting, recipients favor their benefactor over other subjects who have signaled the same intentions but toward others, and this favoritism continues even under a loss of efficiency.

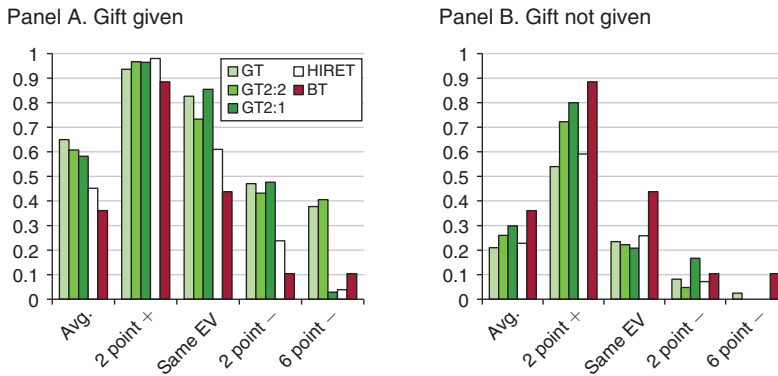


FIGURE 4. RELATIVE FREQUENCIES OF CHOOSING X IN THE CONTROL TREATMENTS COMPARED TO GT AND BT

Passivity of the Client.—So far clients have been passive while producers can make an active choice that benefits DM. This design captures an important feature of many real-world applications. If a doctor prescribes a drug paid for by an insurance company or if a procurement manager acts on behalf of his firm’s shareholders, the client is an anonymous, passive entity. Nevertheless, it is interesting to ask whether DMs behave differently if the client can be more active and also make a choice that benefits DM.

We take a first step toward addressing this question in the *Hiring Treatment* (HIRET). Here, clients can also make a gift of 2 points to DM, at a cost of 1 point to themselves. We frame this as a hiring decision: at the beginning of each period the client can hire DM at a cost of 1 point. If DM is hired, she gets a fixed wage of 20 and chooses the product on behalf of the client as in the gift treatment. If DM is not hired, instead, she gets a fixed outside option payoff of 18, and the product is chosen at random. The client incurs a loss due to the fact that the product is now chosen randomly plus an additional transaction cost of three points. These design features ensure that the total loss of not hiring DM is comparable to the cost of a producer who does not give the gift (and is therefore chosen by DM with a lower probability).

We find that DM is hired in 92.2 percent of all cases. The fourth bars of Figure 4 (panels A and B) display DM’s behavior in HIRET. DMs choose the gift giver’s product in 45.1 percent of all cases (compared to 65.0 percent in GT). Whenever the expected value of X is equal or lower than that of Y , it is chosen much less often than in the GT after gift-giving. Clearly, if the client is in a symmetric position to the producer, the producer’s gift is less effective. This is confirmed by the regression reported in column 6 of Table 3. However, if the gift was not given, DMs’ behavior in HIRET is not significantly different from that in GT, neither on average nor when considering different expected values. Furthermore, column 5 of Table 3 shows that DMs’ behavior in HIRET is still highly significantly different from their behavior in the baseline treatment without gifts. If clients are active, they can reduce but not fully eliminate the effects of gifts given by producers.

Large Externalities.—In the Gift Treatment the externality imposed on the client by DM’s reciprocal behavior does not exceed 6 points in expected value, which is about 40 percent of the client’s expected payoff. To explore whether the propensity to reward gift-giving persists even when the externalities imposed on the client are

TABLE 3—GIFT SIZE AND EFFICIENCY OF GIFT

	Welf.-neutral gift (GT2:2)		Inefficient gift (GT2:1)		Hiring treatment (HIRET)	
	Diff. to BT	Diff. to GT	Diff. to BT	Diff. to GT	Diff. to BT	Diff. to GT
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Model 1: Overall effects</i>						
Respective treatment: gift given	0.241 (0.044)	-0.049 (0.043)	0.216 (0.046)	-0.074 (0.044)	0.090 (0.039)	-0.199 (0.037)
Respective treatment: gift not given	-0.102 (0.054)	0.049 (0.056)	-0.076 (0.054)	0.077 (0.055)	-0.139 (0.048)	0.013 (0.049)
Dummies for (GT: gift given) and (GT: gift not given)	X		X		X	
Dummies for (gift given) and (gift not given)		X		X		X
Controls for gender, econ major, and period	X	X	X	X	X	X
R ²	0.128	0.128	0.123	0.123	0.120	0.120
<i>Model 2: Estimates by EV differences</i>						
(Product X has same EV) × (treatment: gift given)	0.291 (0.078)	-0.098 (0.071)	0.405 (0.066)	0.016 (0.058)	0.170 (0.070)	-0.219 (0.062)
(Product X has same EV) × (treatment: gift not given)	-0.222 (0.090)	-0.018 (0.093)	-0.231 (0.092)	-0.026 (0.095)	-0.183 (0.089)	0.021 (0.092)
(Product X has higher EV [+2]) × (treatment: gift given)	0.075 (0.047)	0.022 (0.040)	0.072 (0.048)	0.019 (0.041)	0.092 (0.039)	0.040 (0.030)
(Product X has higher EV [+2]) × (treatment: gift not given)	-0.166 (0.110)	0.182 (0.127)	-0.095 (0.094)	0.255 (0.113)	-0.298 (0.111)	0.049 (0.127)
(Product X has lower EV [-2]) × (treatment: gift given)	0.325 (0.075)	-0.042 (0.079)	0.369 (0.082)	0.003 (0.085)	0.130 (0.055)	-0.237 (0.060)
(Product X has lower EV [-2]) × (treatment: gift not given)	-0.068 (0.052)	-0.045 (0.054)	0.046 (0.074)	0.070 (0.076)	-0.034 (0.057)	-0.011 (0.058)
(Product X has lower EV [-6]) × (treatment: gift given)	0.297 (0.088)	0.024 (0.091)	-0.083 (0.043)	-0.355 (0.049)	-0.069 (0.043)	-0.342 (0.049)
(Product X has lower EV [-6]) × (treatment: gift not given)	-0.112 (0.033)	-0.037 (0.026)	-0.116 (0.033)	-0.040 (0.026)	-0.105 (0.033)	-0.028 (0.026)
Dummies for EV differences	X	X	X	X	X	X
Dummies for EV differences interacted with (GT: gg) and (GT: gng)	X		X		X	
Dummies for EV differences interacted with (gg) and (gng)		X		X		X
Controls for gender, major, and period	X	X	X	X	X	X
R ²	0.366	0.366	0.392	0.392	0.386	0.386
Sample	GT2:2, GT, BT		GT2:1, GT, BT		HIRET; GT; BT	
Observations	1,680	1,680	1,680	1,680	1,800	1,800

Notes: The HIRET sample contains all observations from the Hiring Treatment, GT2:1 and GT2:2 sample the observations from the Gift Treatments with efficiency 2:1 and 2:2, respectively; BT the observations from the Baseline Treatment; and GT the observations from the Gift Treatment. The abbreviation gg indicates “gift available and given” and gng indicates “gift available but not given.” Constant included. Robust standard errors are reported.

very large, we conducted the *Large Externality Treatment* (LET). Here, we added 10 more periods to 2 sessions of the Gift Treatment. In 6 of these 10 periods the expected value of product X is 16 points lower than that of product Y. For those periods 11 percent of DMs choose X if the gift was given compared to 38 percent if the disadvantage of X was six points and 47 percent if the disadvantage was 2 points.¹⁸ Eleven percent is still a significant fraction, but the responsiveness to the gift is reduced if the externality becomes large.

¹⁸See also the regression reported in columns 7 and 8 of Table 1.

E. Awareness

Are decision makers aware of how strongly their behavior has been influenced by the gift? To answer this question we asked them at the end of the experiment to “estimate in how many periods the product that you chose coincided with the product your client would have chosen by himself.” We also asked them to estimate in how many periods other decision makers had chosen the product the client preferred. Finally, we asked clients and producers the same questions (about the decision makers). All subjects were paid for the precision of their estimates using a quadratic scoring rule.

The result is remarkable: in the Gift Treatment, DMs on average estimate that they chose the client’s preferred product in 69.8 percent of cases; clients predict that DMs chose the preferred product in 68.8 percent, and producers in 68.2 percent of cases. All three estimates are very close to the actual frequency, 65.2 percent. Thus, neither DMs nor clients or producers seem to systematically over- or underestimate the quality of the decisions.¹⁹ However, when decision makers are asked to estimate how often *other* decision makers chose the preferred product of the client, their estimate drops to 63.1 percent. While DMs’ actual behavior is not significantly different, the difference to DMs’ estimates of their *own* behavior is significant at the 1 percent level.²⁰

These results, as well as the finding that DMs match clients’ preferences almost perfectly in the Baseline Treatment imply that DMs are well aware of being influenced by the gift. To better understand this influence, we asked subjects at the end of the experiment several questions about their own motivation and the perceived motivation of the other players. Subjects had to answer these questions by choosing a natural number between 1 (= fully agree) and 6 (= do not agree at all). If the average of the reported numbers is below 3.5, subjects tend to agree with a statement; if the average is above 3.5, they tend to disagree. If a subject reports 1 or 2 (5 or 6), we say that this subject “strongly agrees” (“strongly disagrees”) with the statement.

The first set of questions in the Gift Treatment refers to the motivation of the gift giver. When asked why the producers passed on the gift, almost all decision makers strongly agree with the statement that “the producer wants to influence my behavior” (1.63 on average). DMs are indifferent about the statement that “the producer wants to be nice to me” (3.54) and skeptical about the statement that the producer does so for efficiency reasons because the gift is doubled and “my gain is larger than his” (4.03). Furthermore, they do not agree with the statement that if the producer did not pass on the gift he did so because he “does not want to leave the impression that he wants to influence my decision” (4.75). Clients’ answers to these questions were very similar.

¹⁹ Wilcoxon-Mann-Whitney tests comparing the decisions of DMs with the predictions of the clients or of the producers do not reject the hypothesis that they are drawn from the same distribution ($p = 0.510$ and $p = 0.393$, respectively).

²⁰ A Wilcoxon signed-rank test comparing the decisions of DMs and the predictions of DMs about the behavior of other DMs does not reject the hypothesis that the two are drawn from the same distribution ($p = 0.520$). A Wilcoxon signed rank test comparing the predictions of DMs about their own behavior and about the behavior of other DMs rejects the hypothesis that the two are drawn from the same distribution at $p = 0.001$.

The perceptions of DMs and clients closely match the self-reported motivations of producers. Producers openly admit that they offered the gifts “to influence the decision of the decision maker in my favor” (1.55), and not “to be nice” (3.91) or for efficiency reasons (4.02). They also tend to agree with the statement that “had I not passed on the gift to the decision maker (s)he would not have bought my product” (2.47). In summary, producers pass on the gift because they want to influence the DM’s behavior and because they are afraid that otherwise DMs will not buy their product; DMs and clients perceive this motivation correctly. Nevertheless, DMs respond to the gift.

We also asked DMs directly whether their “decisions have never been influenced by the premium.” About one-third of subjects (31.3 percent) deny any influence. More than one-half (54.2 percent) openly admit that their decisions have been strongly affected. More than one-half also agree strongly with a similar statement about not receiving the gift, “When one of the producers did not pass on the gift even though he could have done so, I did not buy his product.” Furthermore, when asked whether they believe that other DMs have been influenced, they agree even more strongly.

What motivation explains the influence of the gift? In a last set of questions, we elicited DMs’ emotions toward the gift giver. We find that, among those (54.2 percent) who admit to being influenced, 88 percent report positive emotions toward the gift giver or a sense of obligation to buy his product.²¹ At the same time, all of these 88 percent tend to agree with the statement that the gift was given because “the producer wants to influence my behavior,” 91 percent even strongly so.

V. A Model of Social Preferences with Endogenous Reference Groups

The experimental results show a clear pattern of reciprocal behavior. DMs favor producer *X* if he gives the gift, and they discriminate against him if he does not. This contradicts standard theories of outcome-based social preferences such as altruism, maximin preferences, and inequality aversion, as well as type-based models of reciprocity that all predict that DM favors the client. Intention-based reciprocity is consistent with the observed behavior, but also with the opposite. Furthermore, the experiments show that decision makers are fully aware that producers give the gift not because they are kind, but because they want to influence their behavior at the expense of their clients. Thus, gift-giving is interpreted by the decision makers as a selfish act of the producer.

Why, then, is gift-giving so effective? DMs report feeling more positive toward the gift giver and a sense of obligation to reciprocate even though they understand the intentions of the giver. The anthropological and sociological literature is well aware of the fact that gifts create an obligation that is independent of the intention of the gift giver. In a highly influential essay, the anthropologist Mauss (1924) argues that in archaic societies humans are under an “obligation to give, to receive, and then to repay.”²² Prominent sociologists such as Gouldner (1960) and Blau (1964) argue

²¹ When asked whether they “liked a producer who passed on the gift better than the other producer,” 88 percent of DMs tend to agree. When asked whether they “felt obliged to buy the product” of the gift giver, again 65 percent tend to agree, and when asked whether the gift giver “deserves that his product is bought,” 54 percent tend to agree.

²² Mauss (1924) was inspired by Malinowski’s (1922) anthropological field study of the Trobrianders (islanders in the Western Pacific) that identifies 80 forms of social and economic exchange as based on reciprocity.

that reciprocity is a universal social norm that is not just enforced by social pressure and self-interest to maintain a mutually beneficial relationship, but is “internalized.”²³ This view of a gift as creating a bond and an obligation to reciprocate captures the observed behavior better than theories of reciprocity that are based on income distribution, type, or intentions.

In the following we propose a simple extension of outcome-based models of social preferences that formalizes this view. Suppose that, initially, DM is equally concerned about the payoffs of all other players. If another player increases DM’s payoff (and if DM did not perfectly anticipate this), that player gets a higher weight in DM’s utility function. If a player reduces DM’s payoff (as compared to DM’s expectation), the player gets a lower (possibly negative) weight. The key innovation, relative to existing models of action-based reciprocity, is that the weight producer X gets in DM’s utility depends not only on what producer X does, but also on what DM expects him to do.²⁴ If an action was expected with a high probability, the action has less of an effect than if it was expected with a small probability. We endogenize the *reference group* that each player cares about by making the weight on the material payoff of a player in the utility function of another player dependent on the actions of the former and the expectations of the latter.

Formally, consider an N -player game of perfect information. Player $i \in \{1, \dots, N\}$ chooses strategy s_i out of his strategy set S_i . Let $s = (s_1, \dots, s_N)$ denote a pure strategy profile of all players. As in Section III, we assume that all parties are risk neutral. The utility of player i is given by

$$(2) \quad U^i = m^i(s) + \sum_{j \neq i} \alpha_i^j(s | \sigma) \cdot m^j(s),$$

where $\alpha_i^j(s | \sigma)$ is the weight player i puts on the payoff of player j . Thus, i ’s utility depends not only on his own material payoff $m^i(s)$, which is a function of the strategies chosen by all players, but also on the material payoffs of all other players. What is new here is that the weights of these payoffs in i ’s utility function depend on how the strategies chosen compare to the *expected* strategies. The expected strategy profile σ is a (possibly mixed) strategy profile expected to be played in the game, e.g., because of past experience in similar circumstances, or because σ constitutes a social norm, or because σ is an equilibrium of the game that players expect to be played.

ASSUMPTION 1: *A player has social preferences over an endogenously formed reference group that can be represented by (2). If player j chooses a pure strategy s_j*

²³ See also experiments in the social psychology literature (e.g., Whatley et al. 1999). Kolm (2006) argues that people have positive emotions toward a gift giver and feel “moral indebtedness.” Synonyms for saying “thank you” reflect this insight: “much obliged” in English, “je vous suis très obligé” in French, and “ich bin Ihnen sehr verbunden” (literally: “I am bound to you”) in German. The effectiveness of gifts and compliments even if the recipient is aware of ulterior motives, is a building block of “ingratiation” in social psychology (Jones 1964). People comply with requests from those who have done them a favor, even if the favor was unsolicited and if they do not like the gift giver (Regan 1971, Cialdini 1993).

²⁴ The action-based model in Cox, Friedman, and Sadiraj (2008) is different also in that it is restricted to two-stage games with two players and perfect information: cf. Cox, Friedman, and Gjerstad (2007) and the concept of social ties in van Dijk and van Winden (1997).

that increases (decreases) player i 's payoff compared to the (expected) payoff that i would have received if j had chosen the expected strategy σ_j , then the weight of player j 's payoff in player i 's utility increases (decreases) compared to the weight if j had chosen σ_j : $m^i(s_j, \sigma_{-j}) \geq (\leq) m^i(\sigma_j, \sigma_{-j}) \Rightarrow \alpha_i^j(s_j, \sigma_{-j} | \sigma) \geq (\leq) \alpha_i^j(\sigma | \sigma)$.

Let us apply this simple model to our gift-giving game. Suppose that in the Baseline Treatment where no gift can be made, DM puts equal weight α on the client and on producers X and Y , i.e., $\alpha_{DM}^C = \alpha_{DM}^X = \alpha_{DM}^Y = \alpha > 0$. This implies that DM favors the client. Consider now the Gift Treatment and the No Externality Treatment, and suppose DM expects that the gift is given with probability $\sigma_X \in (0,1)$. Thus, if producer X gives the gift, DM's payoff increases to $20 + 2$ as compared to what she expected, $20 + 2\sigma_X$; so the weight that she attaches to the welfare of producer X also increases, $\alpha_{DM}^X(gg | \sigma_X) > \alpha$, where gg indicates "gift given." If producer X does not give the gift, DM's payoff decreases (to 20), compared to what she expected, so the weight that she attaches to producer X also decreases, $\alpha_{DM}^X(gng | \sigma_X) < \alpha$, where gng indicates "gift not given."

In the following, we assume for simplicity that $\alpha_{DM}^X(gg | \sigma_X) = k \cdot \alpha$, with $0 < \alpha < 1$, and where $k > 1$ is distributed across subjects according to some CDF $F(k)$, and $\alpha_{DM}^X(gng | \sigma_X) = 1 \cdot \alpha$, where $0 \leq 1 \leq 1$ is distributed according to some CDF $G(l)$.

PROPOSITION 4: *Consider a decision maker with social preferences satisfying Assumption 1.*

- (i) *Suppose producer X passes on the gift. If product X is (weakly) better than product Y , DM always chooses X in GT and NET. If product X is strictly worse, DM may still choose X , and she is more likely to do so if the payoff goes to a third person (GT) than if it goes to herself (NET).*
- (ii) *Suppose that producer X does not pass on the gift. If X is (weakly) worse than Y , DM always chooses Y in GT and in NET. If X is strictly better, DM may still choose Y , and she is more likely to do so if the payoff goes to a third person (GT) than if it goes to herself (NET).*

PROOF:

See online Appendix A.

The intuition for these results is straightforward. If producer X sends the gift, his weight in DM's utility function increases from α to $k\alpha$. Thus, if X is the better product, DM chooses X in GT and in NET as this choice allows him both to reciprocate toward the gift giver and to benefit her client. If producer X passes on the gift but X is the worse product, DM chooses X in GT if k is large enough to offset the utility loss of the client. In NET DM buys the product for herself, so she has a stronger financial incentive to choose the better product. However, if product X is not much worse than Y , and if k is sufficiently large, DM may still favor producer X since the gift giver gains 16 while the financial cost of reciprocity is small. The argument for case (ii) is analogous.

To transform the framework above into a testable theory we need an *ex ante* measure for the expected behavior. Only given such a measure, we can test whether the weight that the DM puts on the gift giver in her utility function depends on whether and, in the more general version of the model, by how much he exceeded or disappointed her expectations.

One simple proxy for such expectations could be the actual empirical frequency of giving. Remember, for example, that increasing efficiency induces an increase in giving from 63.8 percent in GT2:1 to 67.9 percent in GT2:2 and 73.1 percent in GT. We can make use of this variation to test whether the gift effect depends on how much gift-giving exceeds, or falls short of, expectations by including the relative frequency of gift-giving in a given treatment as one of the independent variables. (Note that, for such an analysis, we return to a more general version of the model, where the weight $\alpha_i^j(s|\sigma)$ that player i puts on the payoff of player j can depend on by how much the j has exceeded or disappointed i 's expectations, rather than fixing those weights at $k\alpha$ and $l\alpha$.)

In Table 4, we estimate the same linear regression model as before but with the added proxy and including data only from treatments where gift-giving is possible, namely, GT2:1, GT2:2, and GT, as well as LGT (which we will discuss in Section VIB). We find that the estimated coefficient for the relative frequency of gift-giving is significantly different from zero and sizable for all specifications. As the model predicts, it is negative, i.e., in a setup where gifts are less likely, reactions to them are stronger. In terms of magnitude, the coefficient is very stable, around -0.28 , regardless of whether we include no additional controls (column 1) or a full slate of additional controls (column 3). In other words, if we raise the expectation of receiving a gift, as proxied by the average frequency of producers sending a gift across a treatment, from 0 to 1, we reduce the response to the gift by about 28 pp. An increase of one standard deviation (which is 9 percent across these treatments) reduces the response to the gift by 2 pp.²⁵

As an alternative proxy for expectations, one may postulate that expectations are most affected by a subject's personal experience of receiving gifts so far. A DM who has received gifts with a high frequency in the past, might have high expectations and thus respond less positively to the next gift and more negatively to not receiving a gift. To test for such dynamic effects in reciprocation and punishment, we reestimate our baseline model on the sample of treatments with gift-giving but include the relative frequency of having received gifts prior to the current round as an explanatory variable.²⁶ (As this cannot be defined in a meaningful way for the first period, we do exclude period 1 in these regressions.) Using this alternative proxy, we have more variation across observations, with a mean experience of 0.753 and a standard deviation of 0.213 for the sample in Table 4. The within-subject variation also allows us to include treatment fixed effects. As the last two columns of Table 4 reveal, we estimate again the expected negative relationship. There is a significant negative effect of past gifts on the positive response to a current gift. The

²⁵Of course, this coefficient captures other treatment differences, too. However, in additional (unreported) regressions that include data from all treatments except BT, the coefficients remain significant and sizably negative, even though the comparison weight of the client changes as he assumes an active role (in HIRET and IT).

²⁶We thank one of our anonymous referees for this suggestion.

TABLE 4—FREQUENCY AND EXPERIENCE OF RECEIVING GIFTS

	Prop. gift giving in treatment			Prop. gifts received (prior period)	
	(1)	(2)	(3)	(4)	(5)
Gift given	0.385 (0.025)	0.389 (0.025)			
Proportion of gift giving in treatment	-0.287 (0.150)	-0.316 (0.152)	-0.328 (0.128)		
Proportion of gifts received in prior periods				-0.138 (0.052)	-0.115 (0.058)
Product <i>X</i> has higher EV [+2]			0.413 (0.061)	0.416 (0.060)	0.418 (0.061)
Product <i>X</i> has lower EV [-2]			-0.140 (0.044)	-0.187 (0.043)	-0.188 (0.043)
Product <i>X</i> has lower EV [-6]			-0.213 (0.038)	-0.216 (0.038)	-0.215 (0.038)
(Product <i>X</i> has same EV) × (gift given)			0.587 (0.041)	0.583 (0.040)	0.589 (0.040)
(Product <i>X</i> has higher EV [+2]) × (gift given)			0.306 (0.052)	0.304 (0.052)	0.307 (0.053)
(Product <i>X</i> has lower EV [-2]) × (gift given)			0.355 (0.037)	0.383 (0.037)	0.390 (0.037)
(Product <i>X</i> has lower EV [-6]) × (gift given)			0.304 (0.033)	0.306 (0.032)	0.309 (0.032)
Controls for gender, major, and period		X	X	X	X
Treatment dummies					X
R^2	0.111	0.114	0.341	0.355	0.359
Sample	GT; LGT; GT2:1; GT2:2	GT; LGT; GT2:1; GT2:2	GT; LGT; GT2:1; GT2:2	GT; LGT; GT2:1; GT2:2	GT; LGT; GT2:1; GT2:2
Observations	1,680	1,680	1,680	1,596	1,596

Notes: Proportion gift giving in treatment denotes the proportion of producers that give the gift if possible, separately for each treatment. Proportion of gifts received in prior periods denotes the proportion of prior rounds in which a given decision maker has received gifts. The LGT sample contains all observations from the Large Gift Treatment; GT2:1 and GT2:2 sample the observations from the Gift Treatments with efficiency 2:1 and 2:2, respectively; and GT the observations from the Gift Treatment. Constant included. Robust standard errors are reported. Sample size for column specifications with proportion of gifts received in prior periods is lower because period 1 is omitted.

coefficient estimates indicate that increase in past gifts by one standard deviation reduces the response to the gift by 3 to 4 pp.

Hence, our conceptual framework appears to offer promise to provide the foundation for a testable model of the formation of reference groups for and strength of social preferences.

VI. Policy Treatments

Most of the remedies proposed in practice to deal with the problem of gift-giving fall in two broad categories: disclosure and size limits. The idea of *disclosure* is to raise awareness about the potential influence of the gift. For example, in the United States and many European countries, political parties and individual politicians have to disclose campaign contributions so that voters can see which interest groups supported them. Another prominent example is the US Physician Payment Sunshine Act (see p. 9503), by which the government requires physicians to report annually

all payments over a cumulative value of \$100 and which, since September 30, 2014, is made available to the public.²⁷ The idea of size limits is to reduce the influence of gift givers. For example, in California, local elected officials may not accept gifts from any single source totaling more than \$420 in a calendar year.²⁸ The Internal Revenue Service does not allow businesses to tax-deduct a business gift costing more than \$25.²⁹ In Minnesota, physicians are not allowed to accept gifts from the pharmaceutical industry that exceed a value of \$50 per year.³⁰ These rules are based on the assumption that large gifts have a large impact on behavior, while the effect of small gifts is small and can be ignored.

We test the effects of such policies in our experimental setup with three additional treatments: the *Disclosure Treatment* (DCT), the *Disclosure with Punishment Treatment* (DCPT) and the *Large Gift Treatment* (LGT). We also conducted an *Incentive Treatment* (IT) in which the client can offer monetary incentives to the decision maker to choose the best product. These additional treatments also provide further tests of our proposed theoretical framework. We conducted four sessions each of DCT and IT and two sessions each of DCPT and LGT. In one of the sessions for DCT, a malfunction in the z-Tree software caused an error in the way treatment was assigned, so that session was excluded from our analysis.

Disclosure and Punishment.—In the DCT, clients observe the producers' and decision maker's actions. The DM knows that the gift and the choice of product are disclosed to the client. There are no other changes relative to the Gift Treatment: the client cannot intervene, reward, or punish DM for her behavior. This design tests whether mere disclosure affects behavior of producers and decision makers, e.g., via raising awareness about the potential influence of the gift or "shaming."

We find that 76.4 percent of producers pass on the gift. The difference to 73.1 percent in GT is not significant. Figure 5 shows the resulting choices of DMs when the gift was given (panel A) and not given (panel B). As before, we display the frequencies for the full sample and for expected value subsamples. Within each set of bars, the leftmost bar shows, for comparison, the GT; the next bar shows the DCT; and the rightmost bar, for comparison, the BT. The remaining bars (in the middle) show the results of the additional treatments (DCPT and LGT), discussed below.

The DCT bars show that, if the gift is sent, DMs choose *X* in 59.6 percent of cases and, if the gift is not sent, in 21.2 percent of those cases. Hence, the average frequencies are close to the GT (65.0 and 20.9 percent, respectively). Also, in the four subgroups, the behavior is mostly similar to the GT. One large difference, though, is that DMs choose *X* only in 22 percent of cases when *X* has a 6 point disadvantage but the gift has been sent, compared to 38 percent in the GT.³¹

²⁷ www.cms.gov/Regulations-and-Guidance/Legislation/National-Physician-Payment-Transparency-Program/Downloads/Affordable-Care-Act-Section-6002-Final-Rule.pdf (accessed November 6, 2016).

²⁸ See <http://www.fppc.ca.gov/> (accessed November 6, 2016).

²⁹ See <http://www.irs.gov/publications/p463/ch03.html> (accessed November 6, 2016).

³⁰ Gardiner Harris, "Minnesota Limits on Gifts to Doctors May Catch On." *New York Times*, October 12, 2007 (<http://www.nytimes.com/2007/10/12/us/12gift.html>).

³¹ Fligner-Policello robust rank-order tests comparing the average decisions in DCT and GT when the gift was (not) given reveal no significant differences between the distributions, overall and controlling for expected value differences.




FIGURE 5. RELATIVE FREQUENCIES OF CHOOSING X IN THE POLICY TREATMENTS COMPARED TO GT AND BT

In Table 5, we evaluate the statistical significance of these effects. Columns 1 and 2 replicate the empirical model of Table 2 for the Disclosure Treatment, pooling the observations from BT, GT, and DCT. The top part of column 1 compares the overall effect of the DCT to the BT, mirroring column 2 of Table 2. The bottom part of column 1 compares the effects of DCT and BT by subsample, mirroring column 5 of Table 2. The coefficient estimates are very similar to those for the GT, both economically and statistically. Column 2 shows the differences of DCT relative to GT (corresponding to columns 3 and 6 in Table 2 for the NET). The economically large reduction in response to a gift when X has a six point lower expected value is significant at the 1 percent level. In fact, when we focus on all cases where Y first-order stochastically dominates X, we estimate a marginally significant reduction of similar magnitude (−0.116, unreported). This suggests that disclosure may be effective when DMs do not have ethical “wobble room.” Still, both regressions confirm that behavior in the DCT is almost identical to the behavior in the original GT without disclosure.³² The questionnaire responses are also very similar across DCT and GT. In DCT, we asked in addition whether DMs agree with the statement: “If the client had not learned whether and by whom I received the gift I would have bought the product of the producer who passed on the gift more often.” We find that they tend to disagree with this statement (4.28), consistent with their observed behavior. These results suggest that disclosing gift givers alone may generally not be an effective way to discipline the behavior of the recipient.

In the second disclosure treatment (DCPT), we added the possibility to punish DM for her choice of product. The client can reduce DM’s payoff by one point, at the cost of one-half point to himself. The information structure is as in DCT: the client is informed about the behavior of X and DM.

The possibility of being punished does not deter producers: 74.3 percent pass on the gift. Moreover, the third bars of Figure 5 reveal that DMs’ choices remain very similar to DCT. DMs choose X in 57.1 percent of all cases when the gift was given (compared to 59.6 and 65.0 percent in the DCT and GT, respectively). When the gift

³²Individual-level clustering and bootstrapping methods yield very similar results. With session-level wild cluster bootstrapping, shown in columns 1 and 2 of online Appendix-Table A6, both the increase after a gift as well as the decrease after not receiving a gift are significantly different from 0, with $p = 0.010$ and $p = 0.090$, respectively.

TABLE 5—POLICY TREATMENTS

	Disclosure (DCT)		Disclosure with punishment (DCPT)			Large gift (LGT)		Incentive-ps (IT-ps)		Incentive-nps (IT-nps)	
	Diff. to BT (1)	Diff. to GT (2)	Diff. to BT (3)	Diff. to GT (4)	Diff. to DCT (5)	Diff. to BT (6)	Diff. to GT (7)	Diff. to BT (8)	Diff. to GT (9)	Diff. to BT (10)	Diff. to GT (11)
<i>Model 1: Overall effects</i>											
Respective treatment:	0.237	-0.052	0.212	-0.077	-0.025	0.159	-0.130	0.246	-0.044	0.388	0.098
gift given	(0.037)	(0.035)	(0.037)	(0.035)	(0.043)	(0.041)	(0.039)	(0.056)	(0.055)	(0.036)	(0.034)
Respective treatment:	-0.144	0.009	-0.087	0.065	0.056	-0.071	0.083	-0.052	0.100	-0.120	0.031
gift not given	(0.050)	(0.051)	(0.051)	(0.053)	(0.064)	(0.109)	(0.109)	(0.074)	(0.074)	(0.047)	(0.048)
Dummies for (GT: gift given) and (GT: gift not given)	X		X		X	X		X		X	
Dummies for (gift given) and (gift not given)		X		X	X		X		X		X
Controls for gender, econ major, and period	X	X	X	X	X	X	X	X	X	X	X
R ²	0.128	0.128	0.12	0.12	0.115	0.116	0.116	0.128	0.128	0.156	0.156
<i>Model 2: Estimates by EV differences</i>											
(Product X has same EV) × (treatment: gift given)	0.301	-0.088	0.289	-0.100	-0.012	0.297	-0.092	0.393	0.004	0.371	-0.019
	(0.063)	(0.055)	(0.065)	(0.057)	(0.070)	(0.069)	(0.061)	(0.081)	(0.074)	(0.064)	(0.055)
(Product X has same EV) × (treatment: gift not given)	-0.107	0.099	-0.191	0.013	-0.087	-0.209	-0.003	-0.199	0.005	-0.191	0.011
	(0.112)	(0.114)	(0.093)	(0.095)	(0.133)	(0.157)	(0.158)	(0.126)	(0.128)	(0.093)	(0.095)
(Product X has higher EV [+2]) × (treatment: gift given)	0.099	0.046	0.072	0.020	-0.026	-0.037	-0.090	0.063	0.010	0.069	0.016
	(0.038)	(0.028)	(0.043)	(0.035)	(0.034)	(0.062)	(0.057)	(0.057)	(0.052)	(0.046)	(0.038)
(Product X has higher EV [+2]) × (treatment: gift not given)	-0.259	0.091	-0.099	0.250	0.158	-0.155	0.197	-0.003	0.346	-0.733	-0.387
	(0.128)	(0.142)	(0.093)	(0.112)	(0.151)	(0.210)	(0.219)	(0.122)	(0.137)	(0.087)	(0.107)
(Product X has lower EV [-2]) × (treatment: gift given)	0.306	-0.060	0.295	-0.072	-0.011	0.199	-0.167	0.223	-0.144	0.505	0.137
	(0.060)	(0.065)	(0.060)	(0.065)	(0.077)	(0.061)	(0.066)	(0.108)	(0.111)	(0.066)	(0.070)
(Product X has lower EV [-2]) × (treatment: gift not given)	-0.097	-0.073	-0.068	-0.044	0.028	0.199	0.224	0.102	0.126	0.088	0.110
	(0.027)	(0.030)	(0.049)	(0.051)	(0.042)	(0.288)	(0.289)	(0.108)	(0.109)	(0.072)	(0.074)
(Product X has lower EV [-6]) × (treatment: gift given)	0.120	-0.152	0.157	-0.116	0.036	0.108	-0.163	0.015	-0.258	0.564	0.290
	(0.068)	(0.072)	(0.067)	(0.071)	(0.084)	(0.070)	(0.074)	(0.084)	(0.088)	(0.078)	(0.081)
(Product X has lower EV [-6]) × (treatment: gift not given)	-0.055	0.022	-0.107	-0.032	-0.053	-0.132	-0.056	-0.099	-0.023	0.270	0.345
	(0.053)	(0.049)	(0.032)	(0.025)	(0.043)	(0.038)	(0.034)	(0.033)	(0.025)	(0.104)	(0.102)
Dummies for EV differences	X	X	X	X	X	X	X	X	X	X	X
Dummies for EV diff. × with (GT: gg) and (GT: gng)	X		X		X	X		X		X	
Dummies for EV diff. × with (gg) and (gng)		X		X	X		X		X		X
Controls for gender, major, and period	X	X	X	X	X	X	X	X	X	X	X
R ²	0.374	0.374	0.369	0.369	0.376	0.363	0.363	0.381	0.381	0.361	0.361
Sample	DCT, GT, BT	DCT, GT, BT	DCPT, GT, BT	DCPT, GT, BT	DCPT, DCT, GT, BT	LGT, GT, BT	LGT, GT, BT	IT-ps, GT, BT	IT-ps, GT, BT	IT-nps, GT, BT	IT-nps, GT, BT
Observations	1,800	1,800	1,800	1,800	2,160	1,680	1,680	1,571	1,571	1,769	1,769

Notes: The DCT, DCPT-FI, and LGT samples contain all data from the Disclosure, Disclosure with Punishment, and Large Gift Treatment, respectively. The data from the Incentive Treatment are split into IT-ps (when the client offered profit sharing) and IT-nps (when the client did not offer profit sharing). The abbreviation gg indicates “gift available and given” and gng indicates “gift available but not given.” Constant included. Robust standard errors are reported.

was not given, 27.6 percent (24.7 and 21.2 percent in DCT and GT, respectively) choose X. It seems that the introduction of a small punishment device in addition to a disclosure does not alter DMs’ behavior much. The regression analysis for DCPT in Table 5 (columns 3–5) confirms this impression. In particular, the regression

comparing DCPT to DCT (column 5) reveals that there is no significant difference; none of the coefficients of the dummy variables containing DCPT are significantly different from zero. One reason for this limited influence of punishment may be that decision makers correctly anticipate that clients will seldom make use of it. Overall, clients in the DCPT punish in only 7.5 percent of all periods.³³

As a caveat, we emphasize that disclosure to peers of a DM may have more powerful shaming effects in practice—at least if there is a strong enough norm of treating clients well. Lacking an environment that reflects the prevalent social norm in a community, as well as for anonymity reasons, public shaming is not readily implementable in a laboratory setting. Still, our laboratory results reveal the limits of disclosure in effectively altering behavior.

Size of the Gift.—In the LGT, we triple the benefit of the gift, keeping its cost constant.³⁴ If producer *X* spends 1 point, DM receives 6 (rather than 2) points, almost one-third of her fixed wage of 20. The LGT allows us to test whether large gifts have a stronger effect on behavior than small gifts.

Because the (relative) price of gift-giving is reduced it is not surprising that producers send the gift more often. In 92.1 percent of cases, potential gift givers pass on the gift, compared to 73.1 percent in the GT. Furthermore, producers strongly confirm in the questionnaire that they passed on the gift to influence DM (average answer of 1.25 as compared to 1.55 in GT), and they strongly negate that the gift has no impact (average answer of 5.13 as compared to 4.52 in GT). Similarly, more clients are strongly convinced that the gift influences the decision maker in the Large Gift Treatment (75 percent) than in the Gift Treatment (52 percent).

Perhaps surprisingly, DMs respond significantly less to the larger gift, as shown in Figure 5 and Table 5 (columns 6–7). When the gift is passed on, DMs choose *X* in 53.4 percent of cases, compared to 65.0 percent in GT.³⁵ The difference in coefficients is highly significant ($p < 0.01$) (top part of column 2).³⁶ If product *X* is worse than *Y*, the effect of the gift is about one-half as strong in LGT as in GT, and is significant at the 5 percent level, both for a 6 point and a 2 point disadvantage. The difference between LGT and GT in the average effect of not passing on the gift is insignificant.

While it may be surprising at first glance that the effect of a large gift is weaker than the effect of a small gift, this result is consistent with our theory in Section V. Recall that the gift is passed on in 92.1 percent of all cases in LGT but only in

³³ We have also conducted a variant where the client has no information about the product (DCPT-PI): in two additional sessions with 48 subjects, shown in online Appendix-Table A7, the treatment resembled DCPT except for the following aspect: the client does not know the identity of the gift giver and which product was chosen and, hence, cannot condition his decision to punish on whether DM was influenced to his advantage or disadvantage. Here clients almost never punish (in only 2.9 percent of cases). Interestingly, DMs choose the gift giver's product more often compared to GT and DCT, both when the gift was given (73.9 percent) and when it was not given (25.3 percent). Overall, the results are not straightforward to interpret and neither support nor contradict our model.

³⁴ In treatments GT2:2 and GT2:1, discussed in Section IVD, we doubled and quadrupled the cost of the gift, relative to the benefit, which had very little effect on the behavior of the producer and DM.

³⁵ A Fligner-Policello robust rank-order test comparing the average decisions of each DM when the gift was given in the GT and the LGT shows that these differences are significant with $p = 0.042$.

³⁶ Individual-level clustering and bootstrapping methods reach very similar results. For example, with session-level wild cluster bootstrapping, the increase after a gift in column 6 (top part) of online Appendix-Table A6 and the difference in increase relative to GT (top part of column 4) are both marginally significant with $p = 0.076$ and $p = 0.099$, respectively.

73.1 percent in GT. Thus, the pure strategy s_j of giving the gift is much closer to the expected strategy σ_j in LGT than in GT, and the model predicts that decision makers will react less to it. Similarly, not giving the gift in LGT is farther from the expected strategy σ_j in LGT than in GT, and therefore the model predicts that decision makers will react more negatively to not receiving the gift in LGT than in GT.³⁷

This finding does not imply that size limits are useless. However, it suggests that size limits on gifts may be insufficient to mitigate the effects of gift-giving. In fact, our results show that small gifts may have an even stronger reciprocal effect than large gifts.

Profit Sharing.—A classical economic approach to counteract the gift effect is to align DM's incentives with the client's payoff. In the Incentive Treatment (IT), we allow the client to give 10 percent of profits to DM (without knowing whether the producer offers a gift or not). He bears one-half of the cost, i.e., loses 5 percent of profits. This cost structure makes the reward symmetric to the gift of producer X who also pays only 50 percent of the gift. Gift and reward are also comparable in size.

In this treatment, 25.4 percent of the clients decided to offer profit sharing. If profit sharing was offered, 67.9 percent of potential gift givers pass on the gift; without profit sharing 67.2 percent do so. The third and fourth bars of Figure 6 illustrate the effect of profit sharing (IT(nps) for cases where the client did not share profit, and IT(ps) for cases where the client did share profits).

As shown in panel A, DMs continue to react positively to the gift of the producer when the client offers profit sharing, choosing X 60.6 percent of the time, but less so than in the GT (65.0 percent). The difference estimate of -0.044 is insignificant as shown in the top portion of Table 5, column 9.³⁸ But the insignificant overall effect hides a large negative reaction when X has 2 and 6 points lower expected value (-14 and -26 pp, with the latter difference being significant at $p < 0.01$; see bottom of Table 5, column 9). After profit sharing, DMs reciprocate to the gift in only 30 or 12 percent of cases when there is a 2 or 6 point disadvantage.

We find the opposite effect if the client does not share profits. In that case, DMs strongly favor the gift giver, even more so than in the GT. For example, when X has a 6 point disadvantage, DMs still chooses X in 66.7 percent of cases when the gift was given! The regressions in Table 5 (lower part of columns 10–11) confirm that both the increases over the BT, estimated to be 56.4 pp, and over the GT, estimated

³⁷It is interesting to note that in the LGT only 41.7 percent of DMs admit to being influenced by the gift (as compared to 54.2 percent in GT), and those who do report being influenced admit to feeling bad about this (average 2.2 as compared to 4.5 in GT). An alternative explanation for the size effect, then, is that the salience of the intention is important. Because a large gift is a more salient indicator that the gift giver wants to influence the behavior of the decision maker (the gift looks more like a bribe), some DMs resist this influence more strongly.

³⁸Fligner-Policello robust rank-order tests comparing DMs' average decisions in IT(ps) and GT when the gift was given (not given) reveal a difference between the treatments that is significant, with $p = 0.040$ ($p = 0.068$). Controlling for the expected-value differences, there are no statistically significant differences. Reestimating the regression model with session-level clustering and bootstrapping methods yield very similar results. For example, with session-level wild cluster bootstrapping, the increase after a gift (when profit is shared) is significant with $p = 0.018$, and the difference in increase relative to the GT is not significant, with $p = 0.688$; see online Appendix-Table A6, columns 8 and 9.




FIGURE 6. RELATIVE FREQUENCIES OF CHOOSING X IN THE INCENTIVE TREATMENT COMPARED TO GT AND BT

to be 29.0 pp, are highly significant.³⁹ The same holds when considering all periods in which X is first-order stochastically.

In summary, if the client offers profit sharing, DMs react positively to the gift but less so than in the GT, and the effect of the gift vanishes when product X is much worse than product Y . This is similar to the No Externality Treatment. However, if the client does not offer profit sharing, the effect of the gift is even stronger than in the GT, in particular when X is much worse than Y . Thus, the possibility of profit sharing can backfire—DMs punish clients for not sharing profits.⁴⁰

These effects are consistent with our theory. As shown in online Appendix A (Proposition 5), our model predicts that the gift effect is moderated if the client shares profits, but it is exacerbated if he does not share. Intuitively, the client strengthens the bond between him and DM by sharing, so his weight in DM's utility function increases. The opposite happens when profits are not shared—the bond is weakened and the weight decreases.⁴¹ This interpretation is also consistent with the questionnaire evidence. The large majority of subjects confirm that they have been influenced by the additional profit sharing of the client. Almost all of them

³⁹Fligner-Policello robust rank-order tests comparing DMs' average decisions in GT and IT(nps) show that these differences are statistically significant if the gift is not given ($p = 0.022$). Controlling for differences in expected values, we find statistically significant differences if the gift is not given and X has either a 2 point advantage ($p = 0.056$) or a 6 point disadvantage ($p = 0.029$). Reestimating the regression model with session-level wild cluster bootstrapping, the increase after a gift (when profit not shared) is significant ($p = 0.019$) while the difference in increase relative to the GT is marginally significant ($p = 0.089$); see online Appendix-Table A6, columns 10–11.

⁴⁰We also conducted two variants of the incentive treatment (with 2 sessions, 48 subjects each), shown in online Appendix-Table A7. In the first treatment (Automatic Incentive Treatment, AIT) 10 percent of the client's profit are added automatically to DM's payoff in each period, at no cost to the client. Here, the gift is given in 61 percent of cases. The gift still has a strong and significant effect, both in the positive and (if not given) in the negative direction, though the positive effect is significantly smaller than in the GT. In the second treatment (Small Incentive Treatment, SIT) the setup is the same except that DM receives the profit share for only 1 out of 20 periods that is determined randomly at the end of the experiment. In this treatment the gift is given in 77 percent of cases, about as often as in the GT (73 percent). DM behavior in this treatment is similar to that in GT when the gift has been given. Surprisingly, when the gift has not been given, DM chooses X more often.

⁴¹Note that DMs punish clients who do not share profits rather harshly when the gift is not given (see Figure 4 panel B, especially if X is the worse product). In the context of our model, the decrease in α associated with a non-sharing client is larger in magnitude than the decrease in α associated with a producer who does not send the gift. This could be a result of how profit sharing and gift-giving are framed. Profit sharing could be perceived as a good (business) practice that clients should undertake, while gift-giving may be seen as bribery and therefore an unfavorable business practice.

(95.7 percent) believe that the client offered the additional reward in order to give the decision maker better incentives to choose the best product for him, and most clients strongly confirm that this is their dominant motive.

The main conclusion from the Incentive Treatment is that rewards aligning the interests of decision maker and client can be highly effective. However, once decision makers become aware that clients could offer profit sharing, they punish clients for not offering the additional reward.

VII. Conclusions

This paper fills a critical gap in the literature on social preferences: the analysis of reciprocal behavior with negative external effects. A person may be “kind” because he wants to influence a decision in his favor at the detriment of a third party. Such situations are common in many industries and other settings (such as lobbying), but the motivating behavioral forces are underexplored.

We have shown that the possibility of gift-giving causes a change in behavior. Gifts induce decision makers to favor the gift giver. If a gift is not given, decision makers discriminate against him, both at the expense of the third party. Gifts are also effective when the decision maker buys the product for herself, but the effect is weaker, in particular when the gift giver’s product is much worse. Standard models of social preferences have difficulties explaining the observed behavior. It is better described by anthropological and sociological theories arguing that gifts create an obligation to reciprocate. A simple extension of standard models of social preferences captures this effect.

How to mitigate the effect of gifts is an important policy issue. Our results show that small gifts can have a large impact, even if they are given unconditionally in a one-shot relationship and if the gift cannot convey any information. In our experimental setup, disclosure and size limits do not reduce the effects of gifts while financial incentives can be highly effective. However, introducing the possibility of financial incentives is a double-edged sword. Once decision makers are aware that additional rewards can be offered, they expect them to be given and punish the client for not doing so. It is also worth noting that the cumulative implementation of several of the nonfinancial policy variations (hiring, disclosure, punishment) might yield a similar effect. A more systematic analysis of possible remedies in the lab and in the field is an important topic for future research.

REFERENCES

- Abbink, Klaus.** 2004. “Staff Rotation as an Anti-Corruption Policy: An Experimental Study.” *European Journal of Political Economy* 20 (4): 887–906.
- Abbink, Klaus, Bernd Irlenbusch, and Elke Renner.** 2002. “An Experimental Bribery Game.” *Journal of Law, Economics, & Organization* 18 (2): 428–54.
- Akerlof, George A.** 1982. “Labor Contracts as Partial Gift Exchange.” *Quarterly Journal of Economics* 97 (4): 543–69.
- Andreoni, James, and John Miller.** 2002. “Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism.” *Econometrica* 70 (2): 737–53.
- Blau, Peter.** 1964. *Exchange and Power in Social Life*. New York: Wiley.
- Blumenthal, David.** 2004. “Doctors and Drug Companies.” *New England Journal of Medicine* 351 (18): 1885–90.

- Bolton, Gary E., and Axel Ockenfels.** 2000. "ERC: A Theory of Equity, Reciprocity, and Competition." *American Economic Review* 90 (1): 166–93.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller.** 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–27.
- Campbell, Eric G., Russell L. Gruen, James Mountford, Lawrence G. Miller, Paul D. Cleary, and David Blumenthal.** 2007. "A National Survey of Physician-Industry Relationships." *New England Journal of Medicine* 356 (17): 1742–50.
- Charness, Gary, and Martin Dufwenberg.** 2006. "Promises and Partnership." *Econometrica* 74 (6): 1579–1601.
- Charness, Gary, and Matthew Rabin.** 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117 (3): 817–69.
- Cialdini, Robert B.** 1993. *Influence: the Psychology of Persuasion*. New York: William Morrow and Company.
- Cox, James C., Daniel Friedman, and Steven Gjerstad.** 2007. "A Tractable Model of Reciprocity and Fairness." *Games and Economic Behavior* 59 (1): 17–45.
- Cox, James C., Daniel Friedman, and Vjollca Sadiraj.** 2008. "Revealed Altruism." *Econometrica* 76 (1): 31–69.
- Currie, Janet, Wanchuan Lin, and Juanjuan Meng.** 2013. "Social Networks and Externalities from Gift Exchange: Evidence from a Field Experiment." *Journal of Public Economics* 107: 19–30.
- Dana, Jason, and George Loewenstein.** 2003. "A Social Science Perspective on Gifts to Physicians from Industry." *Journal of the American Medical Association* 290 (2): 252–55.
- Dufwenberg, Martin, and Georg Kirchsteiger.** 2004. "A Theory of Sequential Reciprocity." *Games and Economic Behavior* 47 (2): 268–98.
- Fehr, Ernst, Lorenz Goette, and Christian Zehnder.** 2009. "A Behavioral Account of the Labor Market: The Role of Fairness Concerns." *Annual Review of Economics* 1 (1): 355–84.
- Fehr, Ernst, Georg Kirchsteiger, and Arno Riedl.** 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *Quarterly Journal of Economics* 108 (2): 437–59.
- Fehr, Ernst, and Klaus M. Schmidt.** 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114 (3): 817–68.
- Fischbacher, Urs.** 2007. "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10 (2): 171–78.
- Gouldner, Alvin W.** 1960. "The Norm of Reciprocity: A Preliminary Statement." *American Sociological Review* 25 (2): 161–78.
- Grande, David.** 2010. "Limiting the Influence of Pharmaceutical Industry Gifts on Physicians: Self-Regulation of Government Intervention?" *Journal of General Internal Medicine* 25 (1): 79–83.
- Greiner, Ben.** 2004. "An Online Recruitment System for Economic Experiments." In *Forschung und wissenschaftliches Rechnen*, edited by Kurt Kremer and Volker Macho, 79–93. Göttingen: GWDG Bericht 63.
- Jones, Edward E.** 1964. *Ingratiation: A Social Psychology Analysis*. New York: Appleton-Century-Croft.
- Katz, Dana, Arthur L. Caplan, and Jon F. Merz.** 2003. "All Gifts Large and Small: Toward an Understanding of the Ethics of Pharmaceutical Industry Gift-Giving." *American Journal of Bioethics* 3 (3): 39–46.
- Kolm, Serge-Christophe.** 2006. "Reciprocity: Its Scope, Rationales, and Consequences." *Handbook of the Economics of Giving, Altruism and Reciprocity* 1: 371–541.
- Levine, David K.** 1998. "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics* 1 (3): 593–622.
- Malinowski, Bronislaw.** 1922. *Argonauts of the Western Pacific*. London: Routledge & Kegan Paul.
- Malmendier, Ulrike, and Klaus M. Schmidt.** 2017. "You Owe Me: Dataset." *American Economic Review*. <https://doi.org/10.1257/aer.20140890>.
- Manacorda, Marco, Edward Miguel, and Andrea Vigorito.** 2011. "Government Transfers and Political Support." *American Economic Journal: Applied Economics* 3 (3): 1–28.
- Mauss, Marcel.** 1924. *The Gift: The Form and Reason for Exchange in Archaic Societies*. London: Routledge & Kegan Paul. (Translated by W. D. Halls. London: Routledge, 1990).
- Morgan, M. A., Jason Dana, George Loewenstein, S. Zinberg, and J. Schulkin.** 2006. "Interactions of Doctors with the Pharmaceutical Industry." *Journal of Medical Ethics* 32 (10): 559–63.
- Murphy, Patrick E.** 1995. "Corporate Ethics Statements: Current Status and Future Prospects." *Journal of Business Ethics* 14 (9): 727–40.
- Pan, Xiaofei, and Erte Xiao.** 2016. "It's Not Just the Thought that Counts: An Experimental Study on Hidden Cost of Giving." *Journal of Public Economics* 138 (1): 22–31.
- Rabin, Matthew.** 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review* 83 (5): 1281–1302.

- Regan, Dennis T.** 1971. "Effects of a Favor and Liking on Compliance." *Journal of Experimental Social Psychology* 7 (6): 627–39.
- Ruffle, Bradley J.** 1999. "Gift Giving with Emotions." *Journal of Economic Behavior & Organization* 39 (4): 399–420.
- Steinman, Michael A., Michael G. Shlipak, and Stephen J. McPhee.** 2001. "Of Principles and Pens: Attitudes and Practices of Medicine Housestaff Toward Pharmaceutical Industry Promotions." *American Journal of Medicine* 110 (7): 551–57.
- Susman, Thomas M.** 2008. "Private Ethics, Public Conduct: An Essay on Ethical Lobbying, Campaign Contributions, Reciprocity, and the Public Good." *Stanford Law & Policy Review* 19 (1): 10–22.
- van Dijk, Frans, and Frans van Winden.** 1997. "Dynamics of Social Ties and Local Public Good Provision." *Journal of Public Economics* 64 (3): 323–41.
- Wazana, Ashley.** 2000. "Physicians and the Pharmaceutical Industry: Is a Gift Ever Just a Gift?" *Journal of the American Medical Association* 283 (3): 373–80.
- Whatley, Mark A., J. Matthew Webster, Richard H. Smith, and Adele Rhodes.** 1999. "The Effect of a Favor on Public and Private Compliance: How Internalized Is the Norm of Reciprocity?" *Basic and Applied Social Psychology* 21 (3): 251–59.