CESifo Working Papers
www.cesifo.org/wp

You Owe Me

Ulrike Malmendier
Klaus M. Schmidt

CESIFO WORKING PAPER NO. 4007
CATEGORY 13: BEHAVIOURAL ECONOMICS
NOVEMBER 2012

CESifo *Working Paper No. 4007*

You Owe Me

Abstract

In many cultures and industries gifts are given in order to influence the recipient, often at the expense of a third party. Examples include business gifts of firms and lobbyists. In a series of experiments, we show that, even without incentive or informational effects, small gifts strongly influence the recipient's behavior in favor of the gift giver, in particular when a third party bears the cost. Subjects are well aware that the gift is given to influence their behavior but reciprocate nevertheless. Withholding the gift triggers a strong negative response. These findings are inconsistent with the most prominent models of social preferences. We propose an extension of existing theories to capture the observed behavior by endogenizing the "reference group" to whom social preferences are applied. We also show that disclosure and size limits are not effective in reducing the effect of gifts, consistent with our model. Financial incentives ameliorate the effect of the gift but backfire when available but not provided.

JEL-Code: C910, D730, I110.

Keywords: gift exchange, externalities, lobbyism, corruption, reciprocity, social preferences.

Ulrike Malmendier
Department of Economics
University of California at Berkeley
501 Evans Hall, Berkeley
USA – CA 94720-3880
ulrike@econ.berkeley.edu

Klaus M. Schmidt
Department of Economics
University of Munich
Ludwigstrasse 28
Germany – 80539 Munich
klaus.schmidt@lmu.de

November 8, 2012
We would like to thank Yan Chen, Rachel Croson, Robert Dur, Ernst Fehr, Rudolf Kerschbamer, Matthias Kräkel, George Loewenstein, Cristina Nistor and 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. Leonie Giessing, Jana Jarecki, Sara Kwasnick, Eddie Ning, and Slobodan Sudaric provided excellent research assistance. Financial support by Deutsche Forschungsgemeinschaft through SFB-TR 15, the Excellence Initiative of the German government, and the Alfred P. Sloan Foundation is gratefully acknowledged.

1 Introduction

Social preferences influence individual behavior in many economically important settings, including bargaining, incentive provision at the workplace, cooperation and competition in small groups, and donations to charity (Camerer, 2003; Fehr and Gächter, 2000; Sobel 2005). A prominent example is gift exchange. Following Akerlof's (1982) seminal paper on gift exchange in labor markets, a large experimental literature has shown that gifts induce cooperation, both in the laboratory and in the field (see, e.g., Fehr, Kirchsteiger and Riedl, 1993; Gneezy and List, 2006; Falk, 2007). We also have a well-developed body of social-preference theories that explain the observed gift-exchange behavior, such as intention-based reciprocity (Rabin, 1993) or inequity aversion (Fehr and Schmidt, 1999).

Much less attention has been devoted to the dark side of pro-social behavior, the presence of negative externalities. Pro-social behavior towards one person may come at the expense of another person. In the case of gift exchange, the purpose of the gift is to influence behavior, often to the detriment of a third party. Consider, for example, the wide-spread 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 precious bottles of wine or event tickets) from a supplier, who hopes to get favorable treatment relative to his competitors, even if his competitors offer better or cheaper products.[1] Similarly, politicians and regulators receive gifts or campaign contributions from lobbyists trying to affect their decisions in favor of special interest groups. In both examples the recipient of the gift makes a decision on behalf of a "client" who is often anonymous: the shareholders of the procurement manager and the general public. Such practices have raised concerns – and stirred a regulatory debate – about the influence of gifts. Gift giving has been blamed as a major contributor to weak corporate governance, to the dramatic rise of health care costs, and to wasteful pork barrel politics.[2] Nevertheless, these externalities have been largely ignored in the theoretical and experimental literature.

In this paper we use a controlled laboratory study to explore why and to what extent gifts

[1] An extreme example is the pharmaceutical industry that has been estimated to spend USD 8,000-15,000 per year on each physician in the US for marketing, including luxurious dinners, conferences at attractive locations, and generous honoraria (Blumenthal, 2004, p. 1885).

[2] See e.g. Katz et al (2003), Blumenthal (2004), Susman (2008). Policy initiatives addressing these problems range from voluntary codes of conduct (see e.g. Murphy [1995] on corporate ethics statements of large U.S. corporations and Grande [2009] on self-regulation in the pharmaceutical industry) to regulatory reforms and laws limiting the possibilities for gift giving and requiring disclosure, such as the Lobbying Disclosure Act of 1995 and the Honest Leadership and Open Government Act of 2007 in the US.

with negative externalities are effective – and what can be done to mitigate their effects. Standard economic theory can explain the effectiveness of gifts in the case of a repeated relationship or if the gift has informational content. For example, a physician may prescribe more drugs of a pharmaceutical company after attending a conference sponsored by that company because he wants to get more sponsoring in the future or because of scientific information provided at the conference. The main contribution of our paper is to show that there is an additional and powerful effect of the gift *per se*. Subjects reciprocate to (small) gifts even if there are no monetary incentives for doing so and if the gift does not convey positive information about the product. The laboratory setting allows us to exclude future interaction, informational content, or any (other) monetary incentives as explanations for such a response. We show a significant effect even for small-scale gifts that amount to little compared to the income of the recipient. We also find that the effect is significantly stronger when it comes at the expense of a third party, compared to the classic gift-exchange situation with two parties.

We consider a situation in which a decision maker has to make a decision on behalf of a client. Before making the decision he may receive a gift from an interested party. We show that existing theories, including theories of social preferences, are inconsistent with the recipient's response to the gift. In fact, the parameters of the experiment have been chosen such that most theories, including altruism, inequality aversion, maximin preferences, and various theories of reciprocity, predict that the recipient does <u>not</u> favor the gift giver at the expense of a third person. Nevertheless, we find that the effect of gifts is large. It is also significantly stronger than in the classic gift-exchange game without a third party. 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. The recipients are fully aware of this intention but reciprocate nevertheless.

In our experiment, a decision maker has to buy one of two possible products on behalf of a client. The products are simple 50/50 lotteries. The decision maker is instructed to choose the product that is best for her client and is paid a fixed wage, independent of her choice. Before she makes the decision she may receive a small gift from one of the two producers. The gift is given unconditionally and before the producers learn the payoffs of their products so that the gift cannot contain any information about the quality of the product. The setting is anonymous, and players are re-matched after each round. Hence, the gift does not provide any monetary incentives to favor the gift giver. Moreover, in the main treatment (*Gift Treatment*), neither client nor producers find out which product the decision maker picks. Nevertheless, gifts strongly affect behavior.

After sending a gift, the gift giver's product is picked twice as often than in a *Baseline Treatment* in which there is no possibility of gift giving. Even if the product of the gift giver is first-order stochastically dominated by the other product (and its expected value is much lower), almost 50 percent of decision makers choose the product of the gift giver, compared to less than 10 percent in the Baseline. The Baseline Treatment also reveals that decision makers have no problem figuring out what the best product for the client is – their choices coincide with the product the clients prefer in 92 percent of all cases. Hence, the distortion of behavior must be due to the gift.

We also find that *not* giving the gift has a strong effect. 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 the product is the better product (higher expected value), not giving the gift reduces the likelihood that it is chosen from more than 90 percent in a setting without gifts to less than 60 percent if a gift could have been given but the potential gift giver chose not to do so.

Our experimental design allows us to test whether decision makers are aware how strongly the gift affects their behavior. This question is much debated in practice. For example, a questionnaire study by Steinman et al. (2001) found that only 39 percent of medical residents believe that gifts from pharmaceutical companies affect their 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. However, as in Steinman et al (2001), they believe that other decision makers are more strongly affected by the gift than they are.

We then compare the results to a situation without third parties. In the *No Externality Treatment*, there is no client; the decision maker buys the product for herself. We test whether the decision maker reciprocates to the same extent as she does when acting for a client. In that case, gift giving would not reduce efficiency but only redistribute income. We find, however, that the effect of the gift becomes significantly smaller. In particular, when the product of the gift giver is much worse than that of his competitor the effect of the gift vanishes completely.

The most prominent economic theories of other-regarding behavior cannot explain the observed phenomena. Outcome-based theories of social preferences (e.g. altruism or inequity aversion) fail to predict an effect of the gift since favoring one producer harms another producer by the same amount and, in addition, a third person (the client). Theories of type-based or intention-based reciprocity assume that actions affect social preferences by signaling the "type" or the "in-

tentions" of the gift giver. However, in our experiments gift giving sends a negative signal and should not be reciprocated according to these theories.

How can we explain the observed behavior? Our questionnaire evidence suggests that a gift triggers an obligation to repay the gift, independently of the intentions of the gift giver and the distributional consequences. The gift seems to create a special bond between the giver and the recipient, in line with a large anthropological literature documenting that gifts create obligations. Similarly, sociologists argue that many forms of social exchange are based on a universal social norm that gifts have to be reciprocated.

In this paper, we propose to extend existing models of social preferences by endogenizing the reference group. The weight that individual *i* attaches to the welfare of individual *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, and the recipient will reciprocate. The more the favorable act exceeds expectation, the stronger the positive response. This simple model is consistent with the observations in our experiments, including the fact that decision makers punish the potential gift giver for not giving the gift.

Finally, we conduct several policy treatments to evaluate how the problem of gift giving can be mitigated. Most remedies that have been proposed, and sometimes been implemented, fall in two broad categories, disclosure and size limits. In our experimental set-up, we find that disclosure on its own has no effect on behavior. When we inform clients which producer is the potential gift giver, whether he does give the gift, and which product the decision maker chooses, decision makers' behavior remains very similar even though they know that everything is disclosed. This finding is consistent with our proposed model. If recipients reciprocate because the gift has created (or strengthened) a bond, then disclosure should have no "shaming" effect.

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 favor the gift giver in 50 percent of all cases, compared to 68 percent before. This may be surprising at first glance, and is contrary to the logic of size limits, but it is consistent with our theory of social preferences with an endogenous reference group: In our set-up decision makers have higher expectations that the producer will send the gift if it benefits them more, and the reward for favorable acts is smaller if they are expected.

In the last policy treatment, the client can offer financial incentives (profit sharing) to align the payoff of the decision maker with his own interest. With profit sharing the effect of the gift is less pronounced than in the Gift Treatment but still slightly stronger than in the No Exter-

nality Treatment. However, clients choose profit sharing in only 30 percent of all cases. Whenever they do not offer profit sharing (but could have done so), the effect of the gift becomes even stronger than in the Gift Treatment. For example, if the product of the gift giver is much worse than the product of his competitor the decision maker chooses it in almost 70 percent of all cases when the client does not share profits. These results also fit our proposed modeling framework. Profit-sharing increases the weight a decision maker puts on the client's payoff, and hence tilts the balance somewhat back in favor of the client. Deciding to not share profits reduces the weight on the client's payoff, and hence further tilts the balance towards the interest of the gift giver.

The rest of the paper is organized as follows: The next section discusses the relation of our paper to the literature. Section 3 describes the experimental design and the different treatments. Section 4 considers the most prominent economic theories of social preferences and shows that none of them predicts that the decision maker favors the gift giver. Section 5 presents our main experimental results and compares the behavior of the decision makers in the Baseline, Gift, and No Externality Treatments. Furthermore, it analyzes whether decision makers are aware of how gifts affect their behavior, and it reports the questionnaire evidence on motives and beliefs. Section 6 discusses how to explain the observed behavior and offers a theoretical framework to model social preferences with endogenous reference groups. Section 7 considers the policy treatments that test how to mitigate the effects of gift giving. Section 8 concludes.

2 Literature

In addition to the papers mentioned in the Introduction (and the anthropological and sociological literature discussed in Section 7), our paper is related to three branches of the economics literature. First, there is a large experimental literature on gift exchange games, starting from Fehr, Kirchsteiger and Riedl (1993). This literature has established reciprocity as an important motive facilitating gift exchange. However, the theoretical explanations offered to explain reciprocity (e.g. altruism, inequality aversion, maximin preferences, type- and intention-based reciprocity) cannot explain reciprocal behavior in our experiment with externalities.

Our evidence on gifts triggering an obligation to "repay," independently of giver's intentions, instead, is consistent with behavior observed in modified trust games, where the trustee can reciprocate only with some exogenous probability (Strassmair 2009). If this probability is high and the trustor expects a return for his initial gift, his intention is more likely to be "selfish". However, the experimental results show that the behavior of trustees is unaffected by the per-

ceived kindness of the trustor. In most of the literature, however, the gift affects only the giver and the receiver; there are no externalities. A notable exception of a gift exchange game with externalities is the "bribery game" in Abbink et al. (2002) and Abbink (2004), where one player can bribe another player to take an action that is beneficial to him but has negative external effects on the "public" (i.e., all other participants in the experiment). The authors show that repeated interaction can sustain a bribery relationship and that the threat of a (probabilistic) penalties and staff rotation significantly reduces corruption. The focus of our paper is very different. We are interested in the effect of gifts that are legally and socially accepted and given in the absence of repeated interaction or any other monetary incentives.

Second, there is a large empirical literature on the effects of business gifts. Much of this literature focuses on the pharmaceutical industry. In a meta-study based on 29 empirical articles, Wazana (2000) concludes that gifts are "associated with increased prescription rates of the sponsor's medication" (p. 373). Campbell et al. (2007) conducted a survey of 3,167 physicians in six 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 on physicians' opinions on whether it is ethical to accept gifts of the pharmaceutical industry and whether these gifts affect prescription behavior. The general conclusion from this literature is that business gifts are widespread and that they are effective. However, as discussed in Dana and Loewenstein (2003), the empirical literature cannot disentangle the causal factors that explain why gifts work.

A third related strand of the literature are field studies and experiments on the effects of gifts when third parties are involved. For example, Falk (2007) collaborated with a charitable organization and sent out different solicitation letters to 10,000 potential donors. He finds that including a small gift increases the frequency of donations by 17 percent, and a large gift by 75 percent. Manacorda, Miguel and Vigorito (2011) estimate the impact of a large anti-poverty program in Uruguay on political support for the government that implemented it. Those households that benefited from the program are 21 to 28 percentage points more likely to favor the current government than those who did not benefit. These studies suggest that reciprocity effects are not restricted to the lab but extend to the field.

3 Experimental Design

Our experimental design aims to capture some key elements of situations where one person has to rely on another person to make a decision on his behalf, and where the decision affects a third

party that has an interest in influencing the decision maker. We focus on the case where the gift is small (such as a pen, a coffee mug, or an invitation for lunch in the B2B context), where it is given unconditionally, and where the parties interact only once. Such small gifts are common in many cultures and industries and, differently from bribes, are often legally and socially accepted.

In the experiment, there are two producers *A* and *B* who want to sell their products to a client. The client has to buy either product *A* or product *B* 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 for her services. If the decision maker chooses product *A* (product *B*, respectively), producer *A* (*B*) receives a positive (quasi-)rent, while the other producer gets 0. Before DM takes her decision one of the producers can pass on a small monetary gift to her. The gift is unconditional, there is no possibility to refuse the gift if it is passed, and there is no repeated interaction – subjects are anonymously re-matched after each round (imperfect strangers design). The client is aware of the possibility that a gift may be given, but he does not know whether a gift was actually passed on, nor does he observe which product was chosen.

We implemented this set-up as follows. A typical session has 24 subjects: 6 decision makers, 6 clients, 6 producers *A*, and 6 producers *B*.[3] There are 20 periods. In each period the decision maker is anonymously matched with a new client and new producers *A* and *B*. At the beginning of each period one of the two producers is selected as the potential gift giver.[4] This producer receives one additional token from the experimenter that he can either keep or pass on to the decision maker in which case it doubles and DM receives two tokens. To simplify the exposition it will be convenient to relabel the producers and their products. In the following we speak of producer *X* (offering product *X*) if a producer is the potential gift giver in a given period and of producer *Y* (offering product *Y*) if he is not the potential gift giver.

The products *X* and *Y* are simple 50/50 lotteries. The payoffs are natural numbers between 3 and 20 and occur with equal probability. For example product *X* might yield a net profit of either 5 or 11 for the client, while product *Y* yields either 3 or 17. (A table with all 20 lotteries is provided in Appendix B.) It is straightforward to compare the lotteries by expected value, variance, or first-order stochastic dominance. We can classify the 20 periods into four categories:

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

[4] We did not allow for the possibility that both producers can make a gift simultaneously, because in this case the two gifts might simply neutralize each other.

- In four periods, the expected value 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 periods, the expected value of lottery *X* is 2 points lower than the expected value of *Y*.
- In four periods, the expected value of lottery *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*.

The use of lotteries allows us to explore two types of scenarios – one in which a product clearly dominates the other product, and one in which at least some measures of comparison differ, giving the decision maker some leeway in picking one or the other as the "better" product.

The producer who sells his product gets a payoff of 16. The other producer gets 0. The producer who is the potential gift giver must decide whether to pass on 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 the decision maker 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.

The decision maker is paid a fixed wage of 20 tokens per period "for taking a decision that is in the best interest of the client." In each period, she first 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 the product she chooses, and she does not learn how the lotteries resolve.

The client does not know who the potential gift giver is and whether the gift was passed on. He observes the two products and is asked which of them he would have chosen if, hypothetically, he could have made the decision himself. He does not observe which product is actually chosen by the decision maker, nor does he observe the outcome of the lottery. At the end of the experiment he is informed only about the sum of his payoffs in all 20 periods.

The rules of the experiment and the information structure are common knowledge between all subjects. In particular, the instructions that are read aloud at the beginning of each session state explicitly that "(t)he client cannot observe the decision of the decision maker. At the end of period 20 the client learns only the sum of all payoffs received during the experiment."

After 20 periods, subjects are asked to answer a questionnaire. In the first part, decision makers are asked to estimate how often their own decision and the decision of the other decision makers coincided with the preferred product of the clients. Similarly, clients and producers are asked to estimate how often the decision makers chose the product that the clients would have

preferred. The answers to these questions are incentivized with a quadratic scoring rule. In the second part, we ask subjects about their motives for their own decisions and their beliefs about the motives of the other subjects. This will be discussed in more detail in Section 5.4.

We compare the results of this *Gift Treatment* to two other treatments, the *Baseline Treatment* and the *No Externality Treatment*. 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 to the decision maker and gifts are never mentioned. This treatment shows whether decision makers choose the products preferred by their clients if nobody can try to influence them. Comparing the Baseline Treatment to the Gift Treatment 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.

In addition, we will compare the results to three policy treatments (disclosure, variation in gift size, and financial incentives), whose design and results are described in Section 7.

Overall, we conducted 15 sessions with 20-24 participants in each session at MELESSA[5] of the University of Munich in 2010 and 2011. Subjects were undergraduate students of various disciplines from the University of Munich and the Technical University of Munich. As shown in the summary statistics (Table 1), we conducted experiments with overall 356 different subjects, generating a data set of 1,980 observations. The vast majority of students were in the typical age range of 20-29, and slightly more than half (54%) were women. About a quarter were students of economics or business studies. Upon arrival at the lab subjects were randomly and anonymously assigned to the different roles. Sessions lasted about one hour. On average, subjects earned about €12.5 ($17.5), which includes a show-up fee of €4 ($5.60).

4 Behavioral Predictions

To guide our empirical analysis, we present a simple theoretical framework, which allows us to derive the predictions of existing theories of social preferences and reciprocity.

[5] MELESSA is the Munich Experimental Laboratory for the Economic and Social Sciences. All experiments were computerized with the software z-Tree (Fischbacher 2007). The recruitment was done with the software ORSEE (Greiner 2004).

Suppose that the decision maker receives a gift from producer *X*. Is she then going to favor the gift giver or will she choose the product that is in the best interest of her client? The payoff of the decision maker is fixed and unaffected by her decision, including the payoff from the gift (since it is given prior to the decision). The traditional model of rational and self-interested behavior does not predict that the decision maker favors the gift giver. She is indifferent, whether or not the gift is given. If there is no third party (the payoffs go to the decision maker), the traditional model predicts that the decision maker chooses the product that has the better payoff, rather than favoring the gift giver. Thus, if we want to explain why the gift systematically affects her decision we have to look for alternative models of human behavior.

In recent years several theories of other-regarding behavior have been proposed. These theories model how a decision maker could be influenced by the payoffs of other players, their types, or their intentions.[6] In this section we show that, in the context of our experiment, the existing theories fail to predict the observed influence of gift giving. To sum up the main insights derived below, outcome-based theories predict unambiguously that decision makers should *not* be influenced by gifts, but should maximize the expected utility of their clients. Type-based and intention-based theories have multiple equilibria that point in opposite directions. In addition, these latter theories make predictions about frequencies and about players' beliefs that apply to all equilibria and can be tested (and refuted). The rest of this section makes these claims precise.

We assume that the decision maker is risk neutral and evaluates products *A* and *B* by their expected values.[7] We say that DM "favors" producer *i* if she chooses product *i* no matter how it compares to product *j*, with $i, j \in \{X, Y\}$ and $i \neq j$. We say that DM favors the client if she chooses the product with the highest expected value for the client, or if both products have the same expected value the product with the smaller variance. We assume that if the decision maker is indifferent she chooses the product that favors the client.[8] Thus, in the standard, self-interest model the decision maker always chooses the product that favors the client.

Suppose first that the decision maker 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 $i \in \{DM, X, Y, C\}$

[6] See Fehr and Schmidt (2006) and Sobel (2005) for surveys of the literature on "social preferences."

[7] Most 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 (see Rabin, 2000) and, at a first approximation, risk neutrality is not restrictive. Note also that in our experiment the decision maker never observes the outcome of the lotteries.

[8] This assumption is imposed frequently in principal-agent models in which the contracting game has multiple equilibria. It is confirmed by the results of the Baseline Treatment.

and U^{DM} is invariant to permutations of (m^X, m^Y, m^C). We consider the three forms of outcome-based social preferences that have received most interest in the literature: (i) Altruism in the form of utilitarianism (e.g. Andreoni and Miller, 2002) assumes that the utility of the decision maker increases with the sum of the payoffs of the other players. (ii) Maximin preferences (e.g. Charness and Rabin, 2002) assume that DM's utility increases with the payoff of the worst-off in the group. And (iii) inequality aversion (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000) assumes that DM dislikes to be worse off and (to a lesser degree) to be better off than the other players. Note that utilitarianism and maximin preferences are special cases of generalized altruism, which assumes that the utility of a player is weakly increasing in the payoffs of all other players. In the set-up of our experiment, most results hold for any form of generalized altruism.[9]

> **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, the decision maker always chooses the product that favors the client.
> (b) In the *Gift Treatment*, if producer *X* did pass on the gift, the decision maker always chooses the product that favors the client.
> (c) In the *Gift Treatment*, if producer *X* did not pass on the gift, the decision maker favors the client if she is altruistic (utilitarian) or inequality averse, but favors producer *Y* if she has maximin preferences.

Proof: See Appendix A.

Proposition 1 implies that DM never favors *X*, no matter whether a gift is given or not. The proof is not difficult and it is instructive to briefly go through the main arguments. In the Baseline Treatment, the decision maker cannot affect the payoff distribution of producers: One gets a payoff of 16, and the other one a payoff of 0, regardless of the product chosen. But she can affect the expected payoff of the client. Moreover, we designed the experiment such that the payoff of the decision maker is always (weakly) higher than the payoff of any other player in any state of the world. Thus, all three outcome-based theories of social preferences predict that the decision maker favors the client. The same argument holds in the Gift Treatment when producer *X* passed on the gift. Again, DM cannot change the payoff distribution of producers, and her payoff (including

[9] Generalized altruism is sufficient to establish parts (a) and (b) of Proposition 1. In part (c) it matters whether DM cares more about the worst off in the group (i.e., the producer who does not make a sale) than about all other players.

the gift) is now strictly higher than that of any other player; so she favors the client. Finally, if in the Gift Treatment producer *X* did not pass on 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. If DM has maximin preferences, she maximizes the payoff of the worst off in the group and favors producer *Y*.[10]

Consider now models of type-dependent preferences such as Levine (1998) and Gul and Pesendorfer (2010). 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 do not care about payoffs of selfish players. Selfish types care only about their own payoffs. The type of a player is private information. It is common knowledge that, for each player $i \in \{DM, X, Y, C\}$, the probability of being kind is $\mu^{DM} = \mu^X = \mu^Y = \mu^C = \mu$, with $0 < \mu < 1$. Let μ_i^j denote the (updated) belief of player *i* about the type of player j, $i, j \in \{DM, X, Y, C\}$, $i \neq j$. Then the expected utility of a kind player is given by

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

where $\alpha > 0$ is the (common) degree to which a kind player *i* cares about the payoff of a kind player j.[11] The utility function of a selfish player *i* simply is $U^{i,selfish} = m^i$.

> **Proposition 2a.** Suppose that the decision maker has type-dependent preferences as described above. Any Perfect Bayesian Equilibrium in which the selfish type keeps the gift with positive probability 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.

The proof of this proposition is straightforward. Let p^{gg} denote the probability that DM chooses product *X* if the gift was given, and p^{gng} the probability that DM chooses *X* if the gift was not given. A selfish producer *X* keeps the gift only if

[10] If DM has maximin preferences that also account for the sum of all payoffs, as in Charness and Rabin (2002), and she puts sufficiently high weight on the payoff of the worst off she favors *Y*; else she favors the client.

[11] We could have also assumed that a kind player *i* cares about the payoff of a selfish player to the degree $\underline{\alpha}$ and about the payoff of a kind player to the degree $\overline{\alpha}$, with $\underline{\alpha} < \overline{\alpha}$ and $(1-\mu)\underline{\alpha} + \mu\overline{\alpha} > 0$. This complicates the exposition but does not change the qualitative results.

$$1 + p^{gng} \cdot 16 \geq p^{gg} \cdot 16$$

Thus, any equilibrium in which the selfish type keeps the gift with positive probability must have $p^{gg} - p^{gng} \leq 1/16$. (In Appendix A we characterize the set of pooling and separating equilibria of this signaling game in Proposition 2b.) Proposition 2a implies that, if $p^{gg} - p^{gng} > 1/16$, then producer *X* cannot signal that he is the kind type by giving the gift because the selfish type will mimic him. Below, we will see that, in our experiment, $p^{gg} - p^{gng} = 0.47 >> 1/16$.

Consider now models of intention-based reciprocity (Rabin, 1993; Dufwenberg and Kirchsteiger 2004). These models use psychological game theory (Geanakoplos, Pearce, and Stacchetti 1988) to capture the idea that players care not only about the action of the other players but also about their intentions. Psychological games with intention-based reciprocity are consistent with many interesting phenomena, but they also tend to have multiple equilibria. For example, if there are two players that can be kind or hostile to each other, it is an equilibrium that both players are kind because they expect the other player to be kind as well; but it is also an equilibrium that both players are hostile because they expect the other player to be hostile.

To apply intention-based reciprocity to our experiment we simplify the strategy space of DM. We allow DM to choose only between action X, i.e. choosing product *X*, and action C, i.e. favoring the client. Because the experiment has a sequential structure we use the notion of "Sequential Reciprocity Equilibrium" of Dufwenberg and Kirchsteiger (2004), but the same result obtains if we apply Rabin's (1993) notion of "Fairness Equilibrium" to the normal form of the sequential game.

> **Proposition 3a.** Suppose that DM and *X* are motivated by intention-based reciprocity. If producer *X* passes on the gift, DM must conclude that *X*'s intention is kind.

Proof: See Appendix A.[12]

Passing on the gift must always be considered a "kind" act because it increases DM's payoff, while keeping the gift is always "unkind" because it reduces DM's payoff. This is independent of the strategies or beliefs of the players. In the experiment we ask the producers about their intentions when giving the gift and we ask the decision makers about their beliefs what these inten-

[12] In Appendix A we also characterize the set of Sequential Reciprocity Equilibrium outcomes in Proposition 3b. We show that there are equilibria in which producers give the gift and equilibria in which they do not. However, in any equilibrium Proposition 3a must hold.

tions are. This allows us to test this prediction.[13]

5 The Effect of Gifts

5.1 Baseline without Gift Giving

Decision makers are instructed to choose the product that is in the best interest of their clients. Before we can study how this decision is affected by gift giving of interested third parties, we establish what happens if there are no gifts. Which products do decision makers choose, and which products do clients prefer?

In Figure 1, we calculate how frequently product *X* is preferred by the client (grey bars) and how frequently it is chosen by the decision maker (black bars), both on average over all periods (first set of bars) and separately for cases when the expected value of lottery *X* is higher than, equal to, or lower than that of lottery *Y* (second to fifth set of bars). These choices will serve as the benchmark for all other treatments, where producer *X* will have the option of giving a gift.

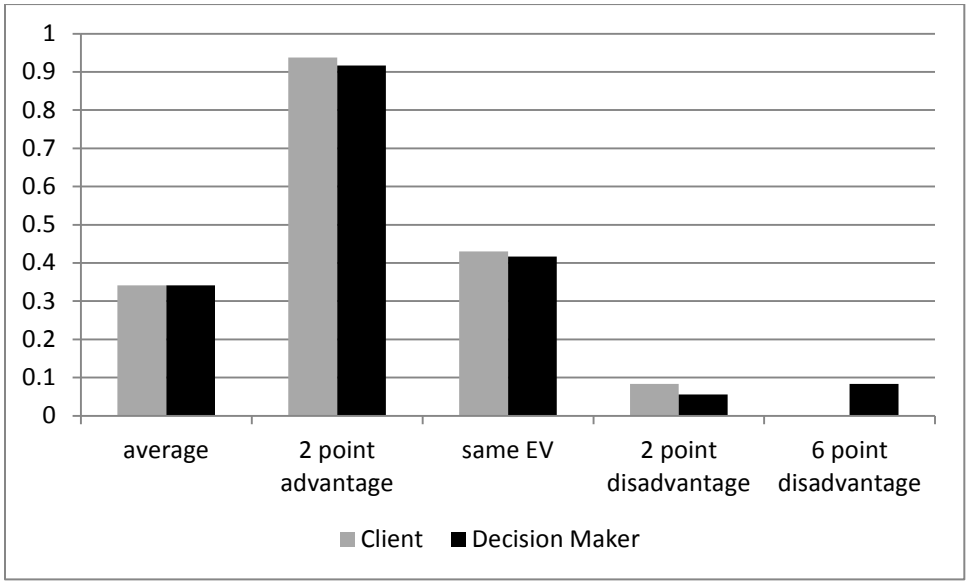

FIGURE 1: Frequencies of choosing *X* in the Baseline Treatment

[13] A related approach is "guilt aversion" (Charness and Dufwenberg, 2006). According to this theory people want to live up to the expectations of others, and they feel guilt if they let other people down. This approach is not easily applicable to our experiment since neither the producer nor the client learn which product DM chooses. They only observe their total payoff at the end of the game. It is difficult to argue that they feel disappointed by DM's decision that they do not observe. Even if the noisy signal that they observe after 20 periods gave rise to disappointment the theory does not offer a clear prediction in our experiment. If product *X* is not the product that is best for the client then the decision maker has to disappoint either the expectations of the gift giver or the expectations of the client. Thus, guilt is unavoidable.

The figure shows that clients strongly prefer the lottery with the higher expected value: 94 percent prefer product *X* if it has the higher expected value. It the expected 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 six points (and is first-order stochastically dominated by *Y*), no client prefers *X*..

The choices of the decision makers closely follow the preferences of the 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.[14] Note that clients prefer and decision makers choose *X* in 34 percent of all cases (82 of 240), not in 50 percent as one might have expected. The reason is that product *X* happens to have the higher expected value in only 4 (out of 20) periods, but a lower expected value in 10 periods. Also note that a little less than 50 percent of clients and decision makers 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 periods. About two thirds of the subjects (70.8 percent of the clients and 55.6 percent of DMs) choose the product with the smaller variance. Hence, in addition to a strong preference for the product with the higher-expected value, we observe a weaker preference for the product with the lower variance.

5.2 Gift Giving with Externalities

What happens if producer *X* can make a gift to the decision maker? Figure 2 illustrates the effect of gifts graphically. The figure compares the choices of the decision makers in the Baseline Treatment (middle bars) to the choices in the Gift Treatment when the gift was passed on (bars to the left) and when it was not passed on (bars to the right). Note that the gift was passed on by the potential gift givers in 71.5 percent of all cases (343 out of 480).

The first set of bars shows a large average effect of the gift: If the gift is given, DMs choose producer *X* twice as often as in the Baseline Treatment, with a frequency of 67.9 rather than 34.2 percent. If, instead, the gift is not given, the frequency decreases by one third, relative to the Baseline Treatment, to 21.2 percent. In other words, decision makers strongly reciprocate to the gift by favoring the gift giver and, perhaps more surprisingly, also exert strong negative reciprocity if a gift could have been passed on but the potential gift giver chose not to do so.

[14] A Wilcoxon-Mann-Whitney test comparing the decisions of DMs and the preferred choices of the clients does not reject the hypothesis that the two are drawn from the same distribution, both when expected values differ and when they are equal ($p = 0.912$ and $p = 0.618$).

The next set of bars shows that, if the lottery offered by the potential gift giver 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 (91.7 and 92.8 percent, respectively). However, when producer *X* does not pass on the gift, the fraction of decision makers buying *X* drops to 55.6 percent. More than 35 percent of DMs "punish" producers for not passing on the gift.

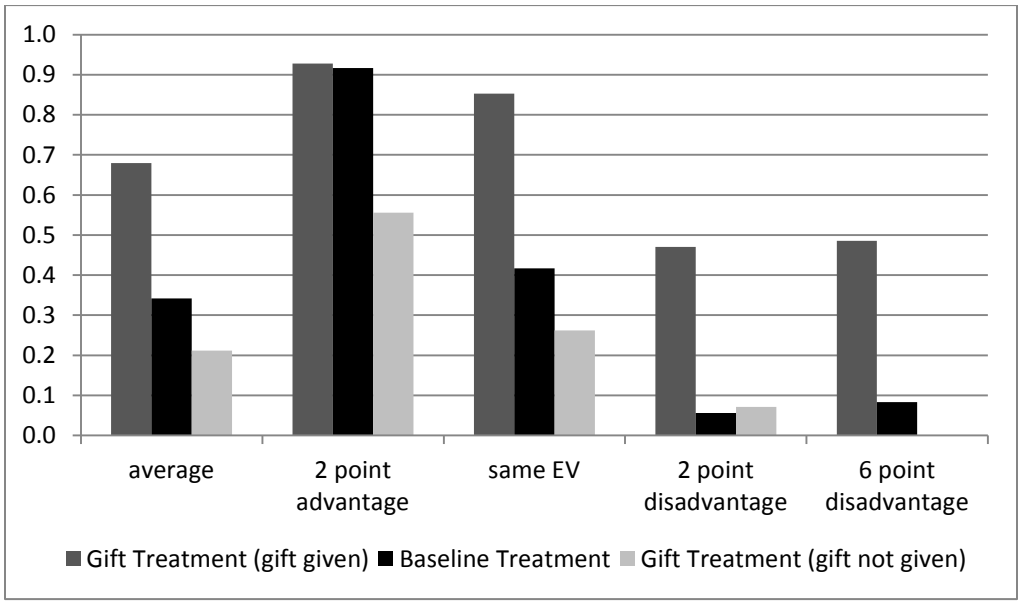

FIGURE 2: Frequencies of choosing *X* in the Gift Treatment (with comparison to the Baseline Treatment)

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* doubles, compared to the Baseline, if the gift was passed on, from 41.7 percent to 85.3 percent, but diminishes by one third, to 26.19 percent, if the gift is not passed on. Finally, if *X* is worse, it is almost never chosen in the Baseline Treatment and, similarly, in the Gift Treatment when the gift was not passed on. However, if producer *X* did pass on the gift, his product is chosen in almost 50 percent of all cases (47.1 percent when the expected value is two points lower, and 48.6 when it is six points lower). Recall, in particular, that if the gift giver's product has a six point disadvantage it is first-order stochastically dominated by product *Y*. Thus there can be no ambiguity what the preferred product of the client is. Nevertheless *X* is chosen almost half of the time when the gift was given.

In Table 2, we evaluate the statistical significance of the effects of gifts in a controlled regression framework. We pool the data of the Baseline and the Gift Treatments and regress an

indicator variable for the choice of product *X* on dummies for *Gift given* and *Gift not given* (in the Gift Treatment), using a linear probability model (OLS) and maximum-likelihood estimation (logit). The regression framework allows us to calculate standard errors clustered by individual and to control for demographics and period-effects..

Columns 1 to 3 show the magnitude and significance of the average effect of the gift, corresponding to the first set of bars in Figure 2. Regardless of the econometric model, giving a gift doubles the probability of the gift giver's product being chosen, and the effect is statistically highly significant ($p < 0.01$). The regressions also confirm that not giving a gift, when gift giving is possible, reduces the probability of *X* being chosen by about one third, and show that the effect is significant at the five-percent level ($p < 0.05$). Adding controls does not affect the coefficient estimates. In particular, neither gender nor the field of study (economics and business administration students versus others) has a significant effect. The *Period* coefficient indicates a slightly negative time trend in decision maker's inclination to choose the gift giver's product, a reduction of about half a percentage point per period, which is marginally significant ($p < 0.10$). However, in unreported results we find that, conditional on receiving or not receiving the gift, there are no significant time trends in the Gift Treatment.[15]

In the next three columns, we refine the regression model to account for the differences in expected values between the two products, corresponding to the second to fifth set of bars in Figure 2. The intercept now captures the case of "no difference in expected value" in the Baseline Treatment. The size and standard error of its coefficient estimate, as well as those of all indicators for expected-value differences ("Product *X* has higher/lower EV …") show that, in a world without gift giving, the influence of product value on product choice is both economically and statistically highly significant: The better product is almost always chosen (with a probability of more than 90 percent), and the worse product is almost never chosen (with a probability of less than 10 percent), and each effect is significant at the one-percent level ($p < 0.01$).

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 coefficient estimates for (EV [+2])*(*Gift given*), (EV [-2])*(*Gift not given*), and (EV [-6])*(*Gift not given*) are cases of left- or right-censoring: Product

[15] When substituting *Period* with three interaction terms of *Period* (with dummies for the Baseline Treatment, *Gift given*, and *Gift not given*), the coefficients of the latter two interactions are insignificant and very small, e.g., in the specification of column 2, -0.005 (s.e. 0.004) and 0.001 (s.e. 0.006) and only the interaction with Baseline Treatment is marginally significant (-0.009 with a s.e. of 0.005).

X is chosen almost always or almost never already in the Baseline Treatment; thus, passing or not passing the gift cannot further alter the probability in a significant manner. Overall, the Gift Treatment shows that gifts can have strong externalities: the client receives the worse product with almost 50 percent probability (instead of less than 10 percent) if the producer of the worse product has sent a gift or if the producer of the better product chose not to send a gift.

5.3 Gift Giving without Externalities

We now ask how the effects of gifts compare, in terms of magnitude and significance, to a setting without externalities. If gifts have the same effect, regardless of who pays the cost, then there is no distortion and it is more difficult to argue that gift giving induces inefficient behavior. If, instead, decision makers behave differently when they act on their own account, then the possibility of gift giving is distortive and likely to be welfare reducing for the third parties.

In the No Externality Treatment there is no client. The decision maker decides on her own behalf and, instead of receiving a fixed wage, is full residual claimant. Figure 3 compares the choices made in the Baseline Treatment (middle bars) to those in the No Externality Treatment when the gift was passed on (bars to the left) and when it was not passed on (bars to the right). Note that 49.1 percent of the potential gift givers (157 out of 320 cases) decide to pass on the gift.

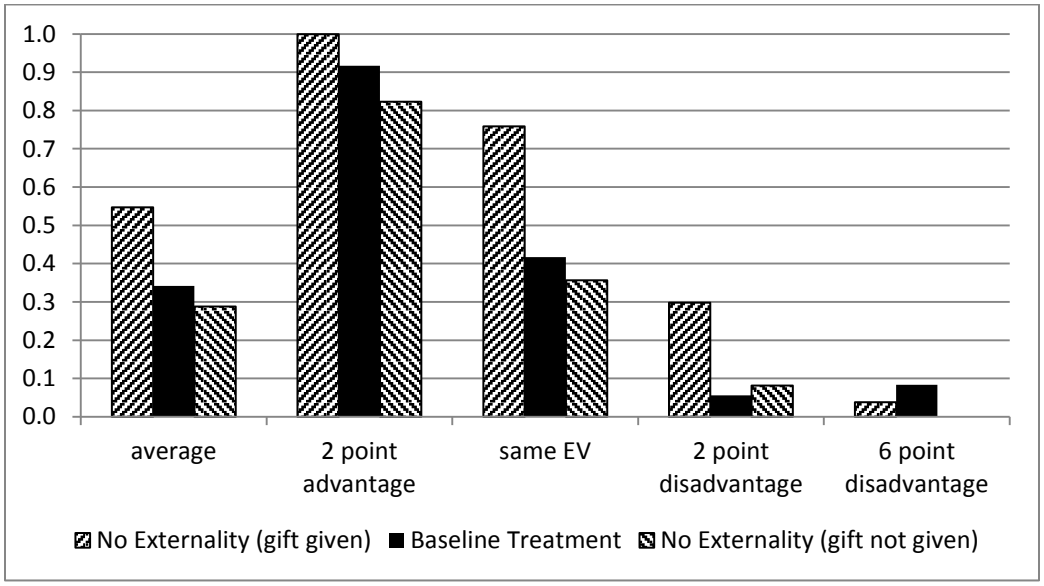

FIGURE 3: Frequencies of choosing *X* in the No Externality Treatment (with comparison to the Baseline Treatment)

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 (34.2 percent in the Baseline Treatment). However, the effects are weaker than in the Gift Treatment, where we observed an increase to 67.9 percent upon receiving the gift and a reduction to 21.2 percent upon not receiving a possible gift.

A closer look at Figure 3 reveals two main differences to Figure 2. First, decision makers punish less often for *not* passing on the gift. Second, when product of the potential gift gives is much worse (6 point disadvantage in expected value) the influence of the gift vanishes.

These findings are confirmed by the regression results, reported in Table 3. For brevity, we only show the linear probability model; the logit results are very similar, both in terms of economics magnitude and in terms of statistical significance. All specifications include the controls for gender, major, and period discussed in the previous table.

The average effects are in columns 1 to 3. We use the pooled data of the Baseline Treatment (BT) and the No Externality Treatment (NET) in column 1, and add the observations from the Gift Treatment (GT) in columns 2 and 3. We find that the increase in the choice of *X* upon receiving a gift, relative to the Baseline Treatment, is smaller but highly significant also in the No Externality Treatment, whether or not we include the data of (and control for) the Gift Treatment (columns 1 and 2). The smaller negative effect of not giving the gift, instead, is insignificant. In other words, we observe less of an increase when the gift is sent and less of a decrease when the gift is not sent when there are no externalities. The difference in increase (about 12.5 percentage points less in the NET than in the GT) is marginally significant, as shown in column 3.

When we control for expected-value differences of *X* and *Y*, we see that the positive effect of receiving the gift is significant in all cases but the expected-value difference of -6 (columns 4). The same is true in column 5 when we add the observations from the Gift Treatment, and control for the corresponding interactions of "*GT: Gift given*" and "*GT: Gift not given.*" Relative to the effect of gifts in the GT, column 6 reveals find significantly less punishment for not sending the gift if *X* is the better product and significantly less reward for sending the gift is *X* is much worse.

There are two main conclusions to be drawn from these results. First, offering a gift is effective even when the decision maker acts on her own behalf. This is consistent with the large experimental evidence on gift exchange games without externalities. It also explains why many firms offer small gifts to their final customers. If the price or quality differences to competing products are not too large, these gifts may tip the balance and induce customers to buy from the gift giving firm. Second, if there are no externalities the effect of the gift disappears when the

product of the gift giver is much worse than the competing product. If product *X* has a six point disadvantage compared to product *Y*, almost 50 percent of all decision makers in the Gift Treatment choose *X* but less than 5 percent in the No Externality Treatment. Thus, the differences in behavior are particularly strong when the external effect is particularly large. This suggests that gift giving can have a large negative impact on social welfare.

5.4 Awareness

Are decision makers aware 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 the other decision makers had chosen the preferred product of the client. Finally we asked the clients and the producers for their estimate of the performance of 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 decision makers on average predict that they chose the preferred product of the client in 64.0 percent of all cases; clients predict that DMs chose the preferred product in 66.4 percent, and producers in 65.3 percent of all cases. All three estimates are very close to each other and to the actual frequency, 63.9 percent. Thus, neither DMs nor clients or producers seem to systematically overestimate or underestimate the quality of the decisions.[16] However, when decision makers are asked to estimate how often *other* decision makers chose the preferred product of the client their estimate drops to 57.9 percent. The difference relative to DMs' actual behavior is significant at the 10 percent level; the difference to DMs' estimates of their *own* behavior is significant at the 5 percent level.[17]

These results, as well as the finding that DMs match clients' preferences almost perfectly in the Baseline Treatment imply that DMs are well aware that they are being influenced by the gift. To better understand this influence, we asked the subjects, at the end of the experiment, several questions about their own motivation and the perceived motivation of the other players. Sub-

[16] A Wilcoxon signed rank test comparing the decisions of DMs and the predictions of their own behavior does not reject the hypothesis that the two are drawn from the same distribution ($p = 0.829$). A Wilcoxon-Mann-Whitney test comparing the decisions of DMs and the predictions of clients and producers does not reject the hypothesis that the two are drawn from the same distribution ($p = 0.764$ and $p = 0.791$).

[17] A Wilcoxon signed rank test comparing the decisions of DMs and the predictions of DMs about the behavior of other DMs rejects the hypothesis that the two are drawn from the same distribution at the 10 percent level ($p = 0.096$). A Wilcoxon signed rank test comparing the predictions of DMs about their own behavior and the behavior of other DMs rejects the hypothesis that the two are drawn from the same distribution at the 5 percent level ($p = 0.024$).

jects 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*" (1.42 on average). Most of them do not agree with the statement that "*The producer wants to be nice to me*" (3.92) or that the producer does so for efficiency reasons because the gift is doubled and "*my gain is larger than his*" (4.71). 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*" (4.67). The answers of the clients to these questions were very similar.

These 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*" (1.82 on average), and not because they wanted to be nice (4.5) or for efficiency reasons (4.67). They also tend to agree with the statement that "(*H*)*ad I not passed on the gift to the decision maker* (*s*)*he would not have bought my product*" (2.59).

In summary, producers report that they pass on the gift because they want to influence the behavior of the decision maker and because they are afraid that otherwise the decision maker will not buy their product, but not because they wanted to be nice or to increase efficiency. Decision makers and clients perceive this motivation correctly. Nevertheless, DMs respond to the gift.

In the next set of questions, we asked DMs directly whether their "*decisions have been influenced*." Only 20.8 percent of the subjects deny any influence. The majority (62.5 percent) openly admit that their decisions have been strongly affected. Furthermore, when asked whether they believe that other DMs have been influenced, they confirm this even more strongly.[18]

What motivation explains the influence of the gift? In a last set of questions, we elicited the emotions of the decision-makers towards the gift giver. We find that, among those (62.5) who admit to having been influenced, 80 percent report positive emotions towards the gift giver or a

[18] 20.8 percent of all DMs strongly disagree (i.e., choose 1 or 2) with the statement "*My decisions have not been influenced*," while 62.5 percent strongly agree (i.e., choose 4 or 5). The average response to this statement is 4.33, compared to an average response of 4.46 to the statement "*The decisions of the other decision makers have not been influenced*." DMs were also confronted with the statement "*When one of the producers did not pass on the gift even though he could have done so, I did not buy his product*." One quarter of the subjects strongly agrees.

sense of obligation to buy his product.[19] At the same time, *all* of these 80 percent also strongly agree with the statement discussed above that the gift was passed on because "The producer wants to influence my behavior."

The survey answers are hard to reconcile with some of the leading explanations for reciprocal behavior, such as intention-based reciprocity.

6 A Framework to Explain the Observed Behavior

The experimental results show a clear pattern of reciprocal behavior. Decision makers favor producer *X* if he gives the gift, and they discriminate against him if he does not. As shown in Section 4, the most prominent theories of social preferences cannot explain these observations. Given the parameters of the experiment, altruism, maximin preferences, and inequality aversion all predict that DM favors the client, not the gift giver. Type-based reciprocity predicts that both the kind and the selfish type will pass on the gift if it increases the probability of being chosen by more than 6.25 percent (and we find that it increases by almost 50 percent); thus we cannot have a separating equilibrium. Intention-based reciprocity implies that giving the gift signals "kind intentions," but decision makers are fully aware that the dominant intention producers want is to influence their behavior to the detriment of their clients.

Why, then, is gift giving so effective? Decision makers report that they feel more positive towards the gift giver and a sense of obligation to reciprocate the gift, even though they understand the intentions of the gift giver. The gift appears to create a special bond between the gift giver and the receiver of the gift.[20] The anthropological, sociological, and social psychology literature is well aware of such an effect, with prominent researchers demonstrating the role of reciprocity as the essential lubricant to sustain social exchange.[21] But this view of reciprocity has

---

[19] When asked whether they "*liked a producer who passed on the gift better than the other producer*," 66.6 percent of DMs strongly agree. When asked whether the gift giver "*deserves that his product is bought*," 53.3 percent strongly agree, and when asked whether they "*felt obliged to buy the product*" of the gift giver, 20.0 percent strongly agree.

[20] This notion is reflected in various synonyms for saying "thank you:" "much obliged" in English, "ich bin Ihnen sehr verbunden" (literally: "I am bound to you") in German, and "je vous suis très obligé" in French.

[21] In a seminal field study of the Trobrianders (islanders in the Western Pacific), the anthropologist Malinowski (1922) identifies 80 forms of social and economic exchange and concludes that all of them are based on reciprocity. In a highly influential essay, Mauss (1924) argues that in archaic societies humans are under an *obligation to give, to receive, and then to repay*. Sociologists such as Gouldner (1960) and Blau (1964) argue that reciprocity is a universal social norm that is not just enforced by social pressure and self-interest to maintain a mutually beneficial relationship in the future, but is often *internalised*. This is confirmed by experiments in social psychology (Whatley et al. 1999). People tend to have positive emotions towards the gift giver and feel a sense of "moral indebtedness" (Kolm 2006) to repay the gift. Many people are willing to 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). The effectiveness of gifts and compliments, even when the recipient is aware of possible ulterior motives, is also a building block in the social

been widely neglected in the economic literature so far.

We suggest a simple model that accounts for the reciprocal behavior in our experiment. The model extends standard outcome-based social preferences by endogenizing the "reference group," or strength of the bond between the parties that are involved. The basic idea is that a gift creates a bond between the giver and the recipient of the gift. Suppose that, initially, DM is equally concerned about the welfare of all other players. Once the gift is given, the welfare of the gift giver gets a higher weight in DM's utility function. Similarly, if a player reduces the decision maker's payoff, that player gets a lower weight.[22]

More formally, consider an *N*-player game of perfect information in which each player $i \in \{1,...,N\}$ chooses strategy s_i out of his strategy set S_i. Let $s = (s_1,...,s_N)$ denote a pure strategy profile of all players. As in Section 4, we assume that all parties are risk neutral. The utility of player i is given by

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

where $\alpha_i^j(s|\sigma)$ is the weight that player i puts on the payoff of player j. Thus, the utility of player i 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 player i's utility function depend on the strategies chosen by these players as compared to the "expected" strategy profile σ. The "expected" strategy profile is a (possibly mixed) strategy profile that players expect to be played in the game under consideration, 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:** If player *j* chooses a pure strategy s_j that increases player *i*'s payoff compared to the (expected) payoff that player *i* would have received if player *j* had chosen the expected strategy σ_j, then the weight of player *j*'s payoff in player *i*'s utility function increases as compared to the weight if *j* had chosen σ, i.e.,

---

psychology literature on "ingratiation," which discusses techniques individuals use to bring themselves into the favor of other people (Jones 1964).

[22] The idea of endogenous reference groups is related to the idea of "social ties" developed by van Dijk and van Winden (1997). They are interested in the dynamics of a repeated public good game and model a social tie as a capital good in the utility function that parties can invest in and that depreciates over time. The idea is also related to Cox et al (2008) who propose a model that may be called "action-based reciprocity." They assume that if a first mover takes an action that increases the maximum attainable payoff of a second mover, then the second mover's preferences will become more altruistic towards the first mover. Their model is restricted to two-stage games with two players and with perfect information and thus not applicable to our set-up, but has the same flavor as our model.

$$m^i(s_j \mid \sigma_{-j}) \geq m^i(\sigma_j \mid \sigma_{-j}) \quad \Rightarrow \quad \alpha_i^j(s_j, \sigma_{-j} \mid \sigma) \geq \alpha_i^j(\sigma \mid \sigma).$$

Let us apply this simple model to our gift giving game. Suppose that in the Baseline Treatment where no gift can be made the decision maker puts equal weight on the client and on producers *X* and *Y*, i.e., $\alpha_{DM}^C = \alpha_{DM}^X = \alpha_{DM}^Y = \alpha > 0$. Therefore, in the Baseline Treatment, 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* passes on 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, $\alpha_{DM}^X(gg \mid \sigma_X) > \alpha$, where *gg* indicates that producer *X* sent the gift ("gift given."). On the other hand, if producer *X* keeps the gift to himself, DM's payoff decreases (to 20) as compared to what she expected, so the weight that she attaches to producer *X* decreases, $\alpha_{DM}^X(gng \mid \sigma_X) < \alpha$, where *gng* indicates that producer *X* did not send the gift ("gift not given.")

In the following proposition, we assume for simplicity that $\alpha_{DM}^X(gg \mid \sigma_X) = \overline{k} \cdot \alpha$, with 0 < α < 1, and where $\overline{k} > 1$ is distributed across subjects according to some cdf $F\left(\overline{k}\right)$. Similarly, we assume that $\alpha_{DM}^X(gng \mid \sigma_X) = \underline{k} \cdot \alpha$, where $0 \leq \underline{k} \leq 1$ is distributed across subjects according to some cdf $G(\underline{k})$. Let $\Delta = m^C(Y) - m^C(X)$ denote the disadvantage of product *X* relative to product *Y* (in terms of expected payoff). As before, we assume that, if indifferent, DM buys the product that favors the client. Then we have:

> **Proposition 4.** Consider a decision maker who has outcome-based social preferences over an endogenously formed reference group.
>
> **(i)** Suppose that producer *X* did pass on the gift. If product *X* is (weakly) better than product *Y*, DM always chooses *X* in the GT and the 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 when it goes to herself (NET).
>
> **(ii)** Suppose that producer *X* did not pass on the gift. If *X* is (weakly) better, DM may choose *X*, but she is less likely to do so when the payoff goes to a third person (GT) than when it goes to herself (NET). If *X* is strictly worse than product *Y*, then DM never chooses *X* in the GT and the NET.

Proof: See Appendix A.

The intuition for these results is straightforward. If producer *X* passes on the gift and *X* is actually the better product, it immediately follows that DM chooses *X* in GT and in NET, as this choice allows DM both to reciprocate towards the gift giver and to benefit her client. If producer *X* passes on the gift but *X* is the worse product, DM might still choose *X* under GT, but only if the increased weight on the gift giver's payoff, from α to $\bar{k} \cdot \alpha$, is large enough to offset the lower product quality for the client. In NET there is no client, and DM has a stronger financial incentive to choose the better lottery. However, if Δ is small, i.e., product *X* is not much worse than product *Y*, and if $\bar{k}$ is sufficiently large, the decision maker 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.

Thus, a simple extension of outcome-based social preferences that endogenizes the reference group is consistent with our experimental observations.

7 Policy Treatments

Several remedies to deal with the problem of gift giving in industry and politics have been proposed and sometimes implemented. Most of these remedies 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 US and many European countries political parties and individual politicians have to disclose campaign contributions so that voters can see which interest groups supported which parties or politicians. Another prominent example is the Physician Payment Sunshine Act in the US under which the government requires annual reporting of all physician payments over a cumulative value of USD 100 and which, after September 30, 2013, will be made available to the public.[23] The idea of size limits is to reduce the maximum influence of gift givers on politicians, physicians and procurement managers. For example, in California local elected officials may not accept gifts from any single source totaling more than US$ 420 in a calendar year.[24] The Internal Revenue Service does not allow businesses to tax-deduct a business gift costing more than US$ 25.[25] In Minnesota physicians are not allowed to accept gifts from the

[23] See http://www.policymed.com/2009/01/physician-payment-sunshine-act-2009-introduced.html, last accessed October 27, 2012.
[24] See http://www.fppc.ca.gov/factsheets/1-09/local_elected.pdf , last accessed October 27, 2012
[25] See http://www.irs.gov/publications/p463/ch03.html , last accessed October 27, 2012.

pharmaceutical industry that exceed a value of US$ 50 per year.[26] 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 these policies in our experimental set-up with two additional treatments, the Disclosure Treatment (DT) and the Large Gift Treatment (LT). We also conducted an Incentive Treatment (IT) in which the client can offer monetary incentives to the decision maker to choose the best product. The additional treatments also provide further tests of our proposed theoretical framework. We conducted two sessions each for the Disclosure Treatment and the Large Gift Treatment and four sessions for the Incentive Treatment.

7.1 Disclosure

In the Disclosure Treatment, clients observe which producer can send a gift, whether this producer does send the gift, and which product is chosen by the decision maker. The decision maker knows that the gift and the choice of product is disclosed to the client. There are no other changes relative to the Gift Treatment: the client cannot intervene, reward, or punish the decision maker for her behavior. This design tests whether mere disclosure affects the behavior of the decision makers, e.g., via raising awareness about the potential influence of the gift or "shaming."

Figure 4 shows the resulting choices of the decision maker, separately for cases where the gift was sent (Panel A) and was not sent (Panel B). As before, we display the overall frequencies and frequencies by subsample, split by differences in expected value. Within each set of bars, the leftmost bar shows, for comparison, the results of the Gift Treatment; the next bar shows the results of the Disclosure Treatment; and the rightmost bar shows, for comparison, the results of the Baseline Treatment. The remaining bars (in the middle) show the results of the other policy treatments and will be discussed below.

The first set of bars in the two panels show that, if the gift is passed on, decision makers choose the gift giver in 65 percent of all cases and, if the gift is not passed on, in 21 percent of all cases in the DT. Hence, the average frequencies are very close to the behavior observed in the Gift Treatment (68 percent and 22 percent, respectively). Also, in the four subgroups of expected-value differences, we mostly observe behavior that is very similar to the Gift Treatment. There is only one large difference: when product *X* has a 6 point disadvantage, decision makers choose the gift giver only in 25 percent of all cases, compared to 50 percent in the Gift Treatment.

[26] See New York Times, October 12, 2007, http://www.nytimes.com/2007/10/12/us/12gift.html , last accessed October 27, 2012.

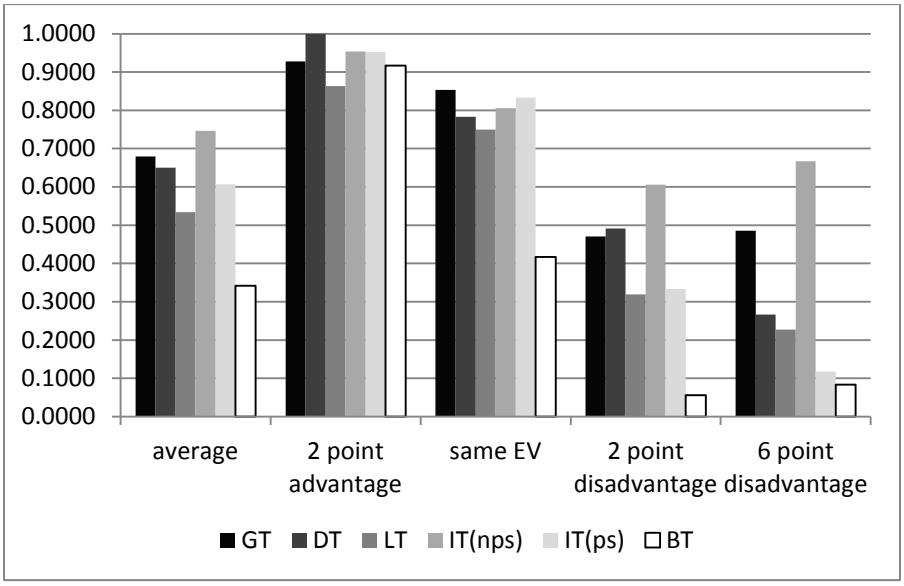

FIGURE 4.A: Frequencies of choosing *X* in the Policy Treatments *when the gift is given* (with comparison to the Baseline and Gift Treatments)

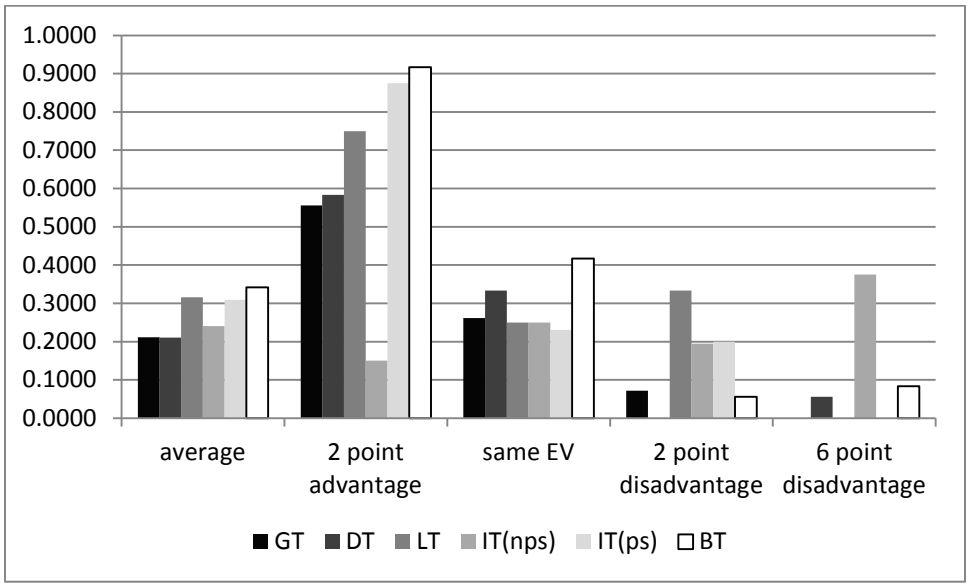

FIGURE 4.B: Frequencies of choosing *X* in the Policy Treatments *when the gift is <u>not</u> given* (with comparison to the Baseline and Gift Treatments)

In Table 4, we evaluate the statistical significance of these effects. Columns 1 and 2 replicate the empirical model of Table 3 for the Disclosure Treatment, pooling the observations from the Baseline, Gift, and Disclosure Treatments. The top part of column 1 mirrors column 1 in Table 3 and compares the overall effect of the DT to the Baseline Treatment. The bottom part of

column 1 corresponds to column 3 in Table 3 and compares the effects of DT and BT by sub-sample. The coefficient estimates are very similar to those in the Gift Treatment, both in terms of economic and in terms of statistical significance. Column 2 shows the differences of the DT relative to the GT (similar to columns 3 and 6 in Table 3 for the NET). None of the differences is significant, including the economically large difference in response to receiving a gift when *X* has a 6 point lower expected value. Hence, both regressions confirm that behavior in the DT is almost identical to the behavior in the original GT without disclosure. In other words, mere "shaming" without actual punishment (such as allowing the client to fire the decision maker) is ineffective in reducing the detrimental effect of the gift for the client.

The answers in the questionnaire are also very similar across DT and GT. In the DT, we asked decision makers in addition whether they 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 (4.17), consistent with their observed behavior. Thus, disclosing the behavior of the decision maker to the client on its own does not seem to be an effective way to discipline decision makers. This finding is consistent with the proposed model. If recipients of a gift reciprocate because the gift has created (or strengthened) a bond between the gift giver and the decision maker, then disclosure should have little effect. That is, re-interpreting reciprocity less in the spirit of the existing economics literature (based on inferring a good "type" or good "intentions) and more in the spirit of the existing sociology, anthropology, and psychology literature, decision makers have little reason to reduce their response when their choices are revealed to others.

7.2 The Size of the Gift

In the Large Gift Treatment (LT) we triple the size of the gift. If the producer passes on the gift the decision maker receives six points rather than two points, which is almost one third of the fixed wage of 20. This treatment allows us to test the hypothesis that large gifts have a stronger effect on behavior than small gifts.

In fact, producers seem to believe that it is important to send the gift in this case. In 92.1 percent of all cases potential gift givers pass on the gift in the Large Gift Treatment compared to 71.4 percent in the Gift Treatment. Furthermore, almost all producers strongly confirm in the questionnaire that they passed on the gift to influence the decision maker (average answer of 1.25 as compared to 1.82 in GT), and they strongly negate that the gift has no impact (average answer

of 5.13 as compared to 4.54 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 (54 percent).

Surprisingly, however, decision makers respond significantly less to the larger gift, as shown in Figures 4.A and 4.B as well as in columns 3 and 4 of Table 4. Consider first the case where the gift is passed on. In this case decision makers choose the gift giver in 50 percent of all cases, compared to 68 percent in the GT. The difference coefficient is significant at the 10 percent level (column 4, top part). If product *X* is worse than *Y*, the effect of the gift is about half as strong in the LT than in the GT, and in the case of a six-point disadvantage the difference is marginally significant. The effect of not passing on the gift is also smaller in the LT than in the GT, though the differences are not significant. (Since the gift was passed on in more than 90 percent of all cases, we have few observations.)

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 6. Recall that the gift is passed on in 92.1 percent of all cases in the LT but only in 71.5 percent in the GT. Thus, the pure strategy s_j of giving the gift is much closer to the expected strategy σ_j in the LT than in the GT, and the model predicts that decision makers will react less to it in LT than in GT.

It is interesting to note that in the LT, only 41.7 percent of the decision makers admit that they have been influenced by the gift (as compared to 62.5 percent in the GT treatment), and those who did report that they have a bad conscience because of this (average 3.3 as compared to 4.3 in GT). An alternative explanation for the size effect, then, is that the salience of the intention has an important impact. Because a large gift is a more salient indicator that the gift giver wants to influence the behavior of the decision maker, some DMs resist this influence more strongly.

This finding has important policy implications. Imposing size limits on gifts is not sufficient to mitigate the effects of gift giving. In fact, small gifts may have a stronger reciprocal effect than large gifts. Thus, not only might size limits be ineffective in reducing the influence of gift giving with externalities, they may even be counterproductive.

7.3 Profit Sharing

The gift induces many decision makers to favor the gift giver. A classical economic approach to counteract this effect is to align the incentives of the decision maker with the payoff of the client. In the Incentive Treatment (IT), we allow the client to give DM 10 percent of his profits (without

knowing whether the producer offers a gift or not). If the client offers profit sharing he bears half of the cost, i.e., loses 5 percent of his profits.[27] In this treatment, 25.4 percent of the clients decided to offer profit sharing. If profit sharing was offered to DMs 67.9 percent of potential gift givers passed on the gift, without profit sharing 67.1 percent did so.

The effect of the option to offer profit sharing is remarkable, as illustrated in Figures 4.A and 4.B, with the fourth and fifth bar ("IT(nps)" for cases where the client decided not to offer a profit share, and "IT(ps)" for observations where the client did offer a share of the profit) in each set of bars. As shown in the first set of bars in Panel A, DMs continue to react positively to the gift of the producer when the client offers profit sharing, choosing *X* 61 percent of the times, but less so than in the Gift Treatment (68 percent). The difference estimate of –0.071 is insignificant as shown in Table 4, column 6. But the insignificant overall effect hides a large negative reaction when *X* has 2 points lower expected value (–14 percent) or 6 points lower expected value (–36 percent), with the latter difference being significant at $p < 0.01$, as shown at the bottom of Table 4, column 6. As a result, decision makers who have been offered profit sharing reciprocate only in about 30 percent of all cases to the gift when there is a 2 point disadvantage, and the effect of the gift completely disappears when there is a 6 point disadvantage.

The opposite is the case if the decision maker has not been offered profit sharing. In that case, DMs strongly favor the gift giver, even more so than in the Gift Treatment. If no profit sharing has been offered and there is a 6 point disadvantage of product *X* decision makers choose the gift giver in 66.7 percent of all cases! The regressions in columns 7 and 8 of Table 4 (lower part) confirm that the increase over the Baseline Treatment, estimated to be 59.2 percentage points, is strongly significant, and the increase relative to the Gift Treatment, estimated to be 19.1 percentage points, is not significant.

In summary, if the client offers profit sharing DMs react positively to the gift but much less so than in the Gift Treatment, 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 Gift Treatment, in particular when product *X* is much worse than product *Y*. Thus, the possibility of profit sharing can backfire – DMs punish clients for not offering profit shares.

These effects are again consistent with our theory in Section 6. In Appendix A (Proposi-

[27] The client has to pay only 50 percent of the cost in order to make his reward symmetric to the gift of producer *X* who also pays only 50 percent of the value of the gift. Gift and reward are also comparable in size.

tion 5), we show that our model predicts that the effect of gift giving is moderated if the client shares profits, similarly to the effect of the gift in the No Externalities treatment. But it is exacerbated if the client does not offer profit sharing. Intuitively, the client strengthens the bond between him and the decision maker by offering profit sharing, so his weight in DM's utility function increases; and the opposite happens when profit sharing is not offered – the bond is weakened and the weight decreases.[28]

This interpretation is also consistent with the questionnaire evidence. The large majority of subjects confirm that their behavior has been influenced by the additional profit sharing offered by the client. Almost all of them (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 that align the interests of the client with the interest of the decision maker can be highly effective. However, once decision makers are aware that clients could offer profit sharing they punish clients if they do not get the additional reward.

8 Conclusions

This paper fills a critical gap in the literature on social preferences: the analysis of situations where reciprocal behavior gives rise to negative external effects. In these situations a person may be "kind" to a decision maker because he wants to influence her decision in his favor at the detriment of a third party. Such situations are common in many industries (business-to-business gifts) and other settings (such as lobbyism), but the motivating behavioral forces are underexplored.

We have shown that the possibility of gift giving causes a change in behavior. If a gift is given the decision maker tends to favor the gift giver; if no gift is given the decision maker tends to discriminate against him, both at the expense of the third party. Gift giving is also effective when the decision maker buys the product for herself, but the effect is much weaker, in particular when product *X* is much worse than product *Y*. This suggests that offering financial incentives

[28] Note that DM appears to punish clients who do not share profits rather harshly when the gift is not given (Figure 4.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. Our model cannot explain why, in the case of profit-sharing (but still no gift giving), an *X* with 2 points disadvantage is chosen 20 percent of the time.

that align the interests of the decision maker and her client may mitigate the negative external effects of gift giving. We show that this is indeed the case. However, introducing the possibility of financial incentives is a two-sided 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.

The most prominent existing theories of social preferences fail to explain the observed behavior. They do not capture the fact that a gift creates an obligation that is largely independent of the intentions of the gift giver and the distributional consequences. One possibility to model this behavior is to allow for reference groups to which social preferences are applied and which form endogenously. By giving or withholding a gift the potential gift giver receives a larger or smaller weight in the utility function of the decision maker.

How to mitigate the effect of business gifts is an important economic 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 gift-giving while financial incentives can be highly effective. A more systematic analysis of possible remedies in the lab and in the field is an important topic for future research.

References

Abbink, Klaus; Irlenbusch, Bernd and Elke Renner. 2002. "An Experimental Bribery Game." *Journal of Law, Economics, and Organization,* Vol. 18, 428–454.

Abbink, Klaus. 2004. "Staff Rotation as an Anti-corruption Policy: An Experimental Study." *European Journal of Political Economy,* Vol. 20, 887–906.

Akerlof, George A. 1982. "Labor Contracts as Partial Gift Exchange." *The Quarterly Journal of Economics*, 97(4): 543–69.

Andreoni, James and John H Miller. 2002. "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism." *Econometrica,* Vol. 70, 737–753.

Blau, Peter. 1964. *Exchange and Power in Social Life*. New York: Wiley.

Blumenthal, David. 2004. "Doctors and Drug Companies." *New England Journal of Medicine,* Vol. 351, 1885–1890.

Bolton, Gary E and Axel Ockenfels. 2000. "ERC: A Theory of Equity, Reciprocity, and Competition." *The American Economic Review,* Vol. 90, 166–193.

Camerer, Colin. 2003. *Behavioral Game Theory. Experiments in Strategic Interaction*. Princeton, N.J.: Princeton University Press.

Campbell, Eric G; Gruen, Russell L; Mountford, James; Miller, Lawrence G; Cleary, Paul D and David Blumenthal. 2007. "A National Survey of Physician-Industry Relationships." *New England Journal of Medicine,* Vol. 356, 1742–1750.

Charness, Gary and Martin Dufwenberg. 2006. "Promises and Partnership." *Econometrica,* Vol. 74, 1579–1601.

Charness, Gary and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *The Quarterly Journal of Economics,* Vol. 117, 817–869.

Cialdini, Robert B. 1993. "Influence: The Psychology of Persuasion." New York: Quill William Morrow.

Cox, James C; Friedman, Daniel and Vjollca Sadiraj. 2008. "Revealed Altruism." *Econometrica,* Vol. 76, 31–69.

Dana, Jason and George Loewenstein. 2003. "A Social Science Perspective on Gifts to Physicians From Industry." *Journal of the American Medical Association,* Vol. 290, 252–255.

Dufwenberg, Martin and Georg Kirchsteiger. 2004. "A Theory of Sequential Reciprocity." *Games and Economic Behavior,* Vol. 47, 268–298.

Falk, Armin. 2007. "Gift Exchange in the Field." *Econometrica,* Vol. 75, 1501–1511.

Fehr, Enst and Simon Gächter. 2000. "Fairness and Retaliation: The Economics of Reciprocity." *Journal of Economic Perspectives*, 14(3): 159–81.

Fehr, Ernst; Kirchsteiger, Georg and Arno Riedl. 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *The Quarterly Journal of Economics,* Vol. 108, 437–459.

Fehr, Ernst and Klaus M Schmidt. 1999. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics,* Vol. 114, 817–868.

Fehr, Ernst and Klaus M Schmidt. 2006. "The Economics of Fairness, Reciprocity and Altruism -

Experimental Evidence and New Theories." In: Kolm, Serge-Christophe; Ythier, Jean Mercier (eds..): *Handbook on the Economics of Giving, Reciprocity and Altruism, Vol. 1*. Amsterdam: Elsevier, 615–691.

Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, Vol. 10, 171–178.

Geanakoplos, John; Pearce, David and Ennio Stacchetti. 1989. "Psychological Games and Sequential Rationality." *Games and Economic Behavior*, Vol. 1, 60–79.

Gneezy, Uri and John List. 2006. "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments." *Econometrica*, 74(5): 1365–84.

Gouldner, Alvin W. 1960. "The Norm of Reciprocity: A Preliminary Statement." *American Sociological Review*, Vol. 25, 161–178.

Grande, David. 2009. "Limiting the Influence of Pharmaceutical Industry Gifts on Physicians: Self Regulation of Government Intervention." *Journal of General Internal Medicine*, Vol. 25, 79–83.

Greiner, Ben. 2004. "An Online Recruitment System for Economic Experiments." In: Kremer, Kurt; Macho, Volker (eds.): *Forschung und wissenschaftliches Rechnen.* Göttingen (GWDG Bericht 63), 79–93.

Gul, Faruk and Wolfgang Pesendorfer. 2010. "The canonical type space for interdependent preferences." Princeton Univerisity,

Jones, Edward E. 1964. *Ingratiation: A social psychological analysis*. New York, New York: Appleton-Century-Croft,

Katz, Dana; Caplan, Arthur L and Jon F Merz. 2003. "All Gifts Large and Small." *The American Journal of Bioethics*, Vol. 3, 39–46.

Levine, David K. 1998. "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics*, Vol. 1, 593–622.

Malinowski, Bronislaw. 1922. *Argonauts of the Western Pacific.* London: Routledge & Kegan Paul.

Manacorda, Marco; Miguel, Edward and Andrea Vigorito. 2009. "Government Transfer and Political Support," *American Economic Journal: Applied Economics*, Vol. 3, 1-28.

Mauss, Marcel. 1990 [1924]. *The Gift. The Form and Reason for Exchange in Archaic Societies.* London: Routledge & Kegan Paul; Routledge.

Morgan, M. A; Dana, Jason and George Loewenstein. 2006. "Interactions of Doctors with the Pharmaceutical Industry." *Journal of Medical Ethics*, Vol. 32, 559–563.

Murphy, Patrick E. 1995. "Corporate ethics statements: Current status and future prospects." *Journal of Business Ethics*, Vol. 14, 727–740.

Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *The American Economic Review*, Vol. 83, 1281–1302.

Rabin, Matthew. 2000. "Risk Aversion and Expected-utility Theory: A Calibration Theorem." *Econometrica*, Vol. 68, 1281–1292.

Regan, Dennis T. 1971. "Effects of a Favor and Liking on Compliance." *Journal of Experimental*

Social Psychology, Vol. 7, 627–639.

Sobel, Joel. 2005. "Interdependent Preferences and Reciprocity." *Journal of Economic Literature*, Vol. 43, 392–436.

Steinman, Michael A; Shlipak, Michael G and Stephen J McPhee. 2001. "Of principles and pens: attitudes and practices of medicine housestaff toward pharmaceutical industry promotions." *The American Journal of Medicine*, Vol. 110, 551–557.

Strassmair, Christina. 2009. "Can Intentions Spoil the Kindness of a Gift? - An Experimental Study," mimeo, University of Munich.

Susman, Thomas M. 2008. "Private Ethics, Public Conduct. An Essay on Ethical Lobbying, Campaign Contributions, Reciprocity, and the Public Good". *Stanford Law and Policy Review*, Vol. 19, 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*, Vol. 283, 373–380.

Whatley, Mark A; Webster, J. M; Smith, Richard H 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*, Vol. 21, 251–259.

Appendix A

Before proving Proposition 1 we have to properly define the three types of social preferences that we consider.

(i) *Altruism* (*utilitarianism*): A decision maker is utilitarian if her utility function is strictly increasing in $\left(m^X + m^Y + m^C\right)$, i.e., $\dfrac{\partial U^{DM}(m^{DM}, m^X, m^Y, m^C)}{\partial(m^X + m^Y + m^C)} > 0$.

(ii) *Maximin preferences*: The decision maker has maximin preferences if

$$U^{DM} = (1-\lambda)\cdot m^{DM} + \lambda \cdot \min\left\{m^{DM}, m^X + m^Y + m^C\right\} > 0,$$

where $0 < \lambda < 1$. Thus, if m^{DM} is unaffected by DM's decision, she maximizes the payoff of the player who is worst off in the group.[29]

(iii) *Inequality aversion*: The decision maker is inequality averse if she wants to minimize the payoff differences between her own payoff and the payoffs of each of the other players (Fehr and Schmidt, 1999), i.e.,

$$U^{DM} = m^{DM} - \frac{\alpha}{3}\sum_j \min\left\{m^j - m^{DM}, 0\right\} - \frac{\beta}{3}\sum_j \min\left\{m^{DM} - m^j, 0\right\} \text{ with } j \in \{X, Y, C\},$$

where $0 < \beta \le \alpha$. An alternative formulation is that she wants her own payoff to be as close as possible to the average payoff of the other players (Bolton and Ockenfels, 2000), i.e.,

$$U^{DM} = U^{DM}\left(m^{DM}, m^{DM} \Big/ \sum_j m^j\right) \text{ with } j \in \{DM, X, Y, C\},$$

where $\dfrac{\partial U^{DM}}{\partial m^j} > 0$, $j \in \{X, Y, C\}$ iff $m^{DM} \Big/ \sum_j m^j > \dfrac{1}{4}$. Note that the material payoff of the decision maker is 20 plus the gift, which is (weakly) greater than the material payoff that any other player can get in any state of the world, and strictly greater than the average payoff of all other players in any state of the world. Thus, like an altruist, an inequality averse decision maker always wants to increase the material payoffs of the other players.

Proof of Proposition 1: We consider the three cases of the proposition in turn. Note that in all cases the decision maker cannot affect her own material payoff.

(a) *Baseline Treatment*. DM cannot affect the distribution of payoffs among the two producers:

[29] Charness and Rabin (2002) consider the case where the decision maker maximizes a weighted sum of her own payoff, the sum of all payoffs, and the payoff of the worst off in the group. This is a convex combination of utilitarianism and maximin preferences. The extension of our results to this case is straightforward.

One of them must get 0, and the other one must get 16. However, the decision does affect the payoff of the client. Since DM's payoff is always weakly greater than the realized payoff of any other player (and strictly greater than the average of the other players' payoffs), all three outcome-based preference models predict that DM maximizes the payoff of the client.

(b) *Gift Treatment, gift given*: As in the Baseline Treatment the distribution of material payoffs of the two producers is unaffected by DM's choice. Thus, she favors the client.

(c) *Gift Treatment, gift not given*: If producer *X* did not pass on the gift, DM can affect the payoff distribution of the producers. If she chooses product *X* the material payoffs of the producers are (17,0), if she chooses *Y* they are (1,16). If she is utilitarian or inequality averse, she is indifferent between these two distributions – in the case of utilitarian preferences because the sum of payoffs is unaffected, in the case of inequality aversion because DM's utility depends on the average difference between her payoff and the payoff of other players' who are behind, regardless of the distribution among those players. Thus, she maximizes the payoff of the client. However, if she has maximin preferences she favors producer *Y* because

$$20 + \min\{20,1,16, m^C(Y)\} > 20 + \min\{20,17,0, m^C(X)\}$$
$$\Leftrightarrow \qquad 1 > 0$$

where $m^C(i) > 0$ is the client's expected payoff if product $i \in \{X, Y\}$ is chosen. ***Q.E.D.***

Proposition 2b. Suppose that the decision maker has type-dependent preferences as described in Section 4.

(a) There always exists a pooling equilibrium in which no gift is given and DM favors the client.

(b) For $1/(8+\alpha+\sqrt{64+\alpha^2}) \le \mu \le 1/16$, there also exists a pure strategy separating equilibrium in which the "kind" type of producer *X* passes on the gift and the "selfish" type keeps it. In this equilibrium the kind type of DM favors *X* if the gift is passed on and favors *Y* is the gift is kept.

Proof of Proposition 2b: (a) We construct a pooling equilibrium in which no gift is given and DM favors the client. If DM is selfish, she is indifferent and will favor the client, as assumed above, no matter whether the gift was passed on or not. If DM is kind and the gift was not passed on, she maximizes the sum of payoffs since, in a pooling equilibrium, she does not learn anything about the types of the other players on the equilibrium path. Hence, a kind DM also favors the

client. Whether it is optimal for producer *X* to keep the gift depends on DM's beliefs about producer *X*'s type if the gift is passed on. Note that this is an out off equilibrium event that happens with probability 0 in equilibrium. Thus, we are free to specify DM's beliefs in this case. If DM believes that if producer *X* passes on the gift he must be a selfish type with probability 1, then, a kind DM favors *Y* because

$$22 + \alpha \cdot \mu \cdot [16 + m^C(Y)] > 22 + \alpha \cdot \mu \cdot [0 + m^C(X)],$$

which is equivalent to

$$16 + m^C(Y) - m^C(X) > 0,$$

always holds since $m^C(Y) - m^C(X) \geq -2$ in all periods. Thus, producer *X* has no incentive to pass on the gift. This proves the existence of the pooling equilibrium.

(b) Now we construct a separating equilibrium in which the gift is passed on if and only if producer *X* is the kind type, and in which the kind type of DM buys product *X* if and only if the gift is given.

Stage 2: The selfish type of DM will always favor the client, whether the gift was given or not. The decision of the kind DM depends on her beliefs about the type of producer *X*. In the separating equilibrium we are considering the kind type of producer *X* passes on the gift while the selfish type keeps it. Thus, if DM observes that the gift was given she believes that producer *X* is kind with probability 1. In this case, it is optimal for the kind DM to choose product *X* if and only if

$$22 + \alpha \cdot 16 + \mu \cdot \alpha \cdot m^C(X) \geq 22 + \mu \cdot \alpha \cdot [16 + m^C(Y)],$$

which is equivalent to

$$\mu \leq \frac{16}{16 + m^C(Y) - m^C(X)}.$$

Note that the difference between the expected value of product *Y* and product *X* is at most 6. Thus, if $\mu \leq \frac{8}{11}$, this condition is always satisfied.

If DM observes instead that producer *X* kept the gift she believes that she faces the selfish type with probability 1. In this case, it is optimal for the kind type of DM to choose product *Y* if and only if

$$20 + \mu \cdot \alpha \cdot (16 + m^C(Y)) \geq 20 + \mu \cdot \alpha \cdot m^C(X)$$

which is always true because $m^C(X) - m^C(Y) \leq 2$ in all periods.

Stage 1: Consider now the decision of producer *X* whether to pass on the gift. Let p^{gg} denote the probability that DM chooses *X* if the gift is given, and p^{gng} the probability that DM chooses *X* when the gift is not given. For a selfish type of producer *X* it is optimal not to pass on the gift if and only if

$$1+16\cdot p^{gng} \geq 16\cdot p^{gg} \quad \Leftrightarrow \quad p^{gg} - p^{gng} \leq \frac{1}{16}$$

Note that if the gift is given the DM believes that she faces the kind type of producer with probability one, so $p^{gg} = \mu\cdot 1+(1-\mu)\cdot p$ where *p* is the probability with which a selfish type of DM chooses product *X*. If the gift is not given DM believes that she faces the kind type with probability zero, so $p^{gng} = \mu\cdot 0+(1-\mu)\cdot p$. Thus, $p^{gg} - p^{gng} = \mu$. Hence, for a selfish type of producer *X* it is optimal not to pass on the gift if and only if

$$\mu \leq 1/16.$$

For a kind type of producer *X* it is optimal to pass on the gift if and only if

$$\begin{aligned}
&p^{gg}\cdot 16+\mu\cdot\alpha\cdot 22+\mu\cdot\alpha\cdot(1-p^{gg})\cdot 16+\mu\cdot\alpha\cdot(p^{gg}\cdot E[m^C(X)]+(1-p^{gg})\cdot E[m^C(Y)])\\
&\geq 1+p^{gng}\cdot 16+\mu\cdot\alpha\cdot 20+\mu\cdot\alpha\cdot(1-p^{gng})\cdot 16+\mu\cdot\alpha\cdot(p^{gng}\cdot E[m^C(X)]+(1-p^{gng})\cdot E[m^C(Y)])\\
&\Leftrightarrow\\
&p^{gg}\cdot 16+\mu\cdot\alpha\cdot 22+\mu\cdot\alpha\cdot(1-p^{gg})\cdot 16+\mu\cdot\alpha\cdot(p^{gg}\cdot E[m^C(X)]-p^{gg}\cdot E[m^C(Y)])\\
&\geq 1+p^{gng}\cdot 16+\mu\cdot\alpha\cdot 20+\mu\cdot\alpha\cdot(1-p^{gng})\cdot 16+\mu\cdot\alpha\cdot(p^{gng}\cdot E[m^C(X)]-p^{gng}\cdot E[m^C(Y)])\\
&\Leftrightarrow\\
&16\cdot(p^{gg}-p^{gng})+2\mu\alpha-16\mu\alpha\cdot(p^{gg}-p^{gng})+\mu\alpha(p^{gg}-p^{gng})(E[m^C(X)]-E[m^C(Y)])\geq 1.
\end{aligned}$$

Using $p^{gg} - p^{gng} = \mu$, we simplify the inequality to

$$16\mu+2\mu\alpha-16\mu^2\alpha+\mu^2\alpha(E[m^C(X)]-E[m^C(Y)])\geq 1.$$

Since producer *X* does not know products *X* and *Y* when deciding on the gift (and assuming that he does not interpret the fact that he was chosen randomly as the potential gift giver as a signal about the quality of his product) we have

$$E[m^C(X)] = E[m^C(Y)].$$

Thus, it is optimal for the kind type of producer *X* to pass on the gift if and only if

$$16\mu+2\mu\alpha-16\mu^2\alpha\geq 1 \Leftrightarrow \mu^2-\frac{8+\alpha}{8\alpha}\mu+\frac{1}{16\alpha}\leq 0.$$

Solving for the points where the above weak inequality is equal to 0 we get

$$\mu_{1,2} = \frac{8+\alpha}{16\alpha} \pm \sqrt{\left(\frac{8+\alpha}{16\alpha}\right)^2 - \frac{1}{16\alpha}} = \frac{8+\alpha}{16\alpha} \pm \sqrt{\frac{64+16\alpha+\alpha^2-16\alpha}{(16\alpha)^2}} = \frac{8+\alpha}{16\alpha} \pm \sqrt{\frac{64+\alpha^2}{(16\alpha)^2}}\,.$$

We can further simplify

$$\mu_1 = \frac{1}{16\alpha}\left((8+\alpha) - \sqrt{64+\alpha^2}\right) = \frac{1}{16\alpha\left((8+\alpha)+\sqrt{64+\alpha^2}\right)}\left((8+\alpha) - \sqrt{64+\alpha^2}\right)\left((8+\alpha)+\sqrt{64+\alpha^2}\right)$$

$$= \frac{1}{16\alpha\left((8+\alpha)+\sqrt{64+\alpha^2}\right)}\left((8+\alpha)^2 - (64+\alpha^2)\right) = \frac{16\alpha}{16\alpha\left((8+\alpha)+\sqrt{64+\alpha^2}\right)} = \frac{1}{8+\alpha+\sqrt{64+\alpha^2}}$$

and, similarly, $\mu_2 = \frac{1}{8+\alpha-\sqrt{64+\alpha^2}}$.

To summarize, a separating equilibrium requires for the kind DM to favor producer *X* after receiving the gift that $\mu \le \frac{8}{11}$, for the selfish producer *X* not to pass on the gift that $\mu \le \frac{1}{16}$ (which implies $\mu \le \frac{8}{11}$), and for the kind producer *X* to pass on the gift that

$$\frac{1}{8+\alpha+\sqrt{64+\alpha^2}} \le \mu \le \frac{1}{8+\alpha-\sqrt{64+\alpha^2}}$$

Furthermore, $\frac{1}{8+\alpha+\sqrt{64+\alpha^2}}$ is always strictly smaller than 1/16. Thus, a pure strategy separating equilibrium exists if and only if $\frac{1}{8+\alpha+\sqrt{64+\alpha^2}} \le \mu \le \frac{1}{16}$, as claimed in the proposition.

Q.E.D.

> **Proposition 3b.** Suppose that DM and *X* are motivated by intention-based reciprocity as described in Section 3. If both producer *X* and DM care strongly enough about the kindness of the intentions of the other player, then there are multiple sequential reciprocity equilibria (SRE). In particular, there exists a SRE in which producer *X* passes on the gift and DM chooses *X*. But there also exists a SRE in which player *X* does not pass on the gift and DM chooses C.

Proof of Propositions 3a and 3b: At stage 1 producer *X* can choose whether to pass on the gift (G) or not to pass on the gift (N). Then DM decides whether to choose *X*'s product (X) or the product that yields the highest expected payoff for the client (C). Since the role of the gift giver was randomly allocated, we maintain that the parties assign an probability of 50 percent of having the better product. The expected payoffs are given in the normal form of this sequential game:

DM \ X	XX	XC	CX	CC
G	16, 22	16, 22	8,22	8, 22
N	17, 20	9, 20	17,20	9, 20

For players $i \in \{X, DM\}$, let a_i denote player *i*'s strategy, b_{ij} player *i*'s belief about the strategy chosen by player *j* (first-order belief), and c_{iji} player *i*'s belief what player *j* believes about *i*'s own strategy (second-order belief), with $j \neq i$. Player i's expected utility is given by

$$U^i(a_i, b_{ij}, c_{iji}) = m^i(a_i, b_{ij}) + n_i \cdot \kappa_{ij}(a_i, b_{ij}) \cdot \lambda_{iji}(b_{ij}, c_{iji})$$

The first term is *i*'s expected monetary payoff. The second term is *i*'s reciprocity payoff. Here, the parameter $n_i \geq 0$ reflects how much *i* cares about the perceived kindness of player *j*, $\lambda_{iji}(b_{ij}, c_{iji})$. The kindness of player i is given by the function

$$\kappa_{ij}(a_i, b_{ij}) = m^j(a_i, b_{ij}) - m^j_{e_i}(b_{ij}).$$

This is the payoff that player i "gives" to j by choosing a_i assuming that *j* chooses b_{ij}, minus the "equitable" payoff of *j* which is defined as the average of the maximum and the minimum payoff that player *i* can "give" to player *j* (assuming that *j* chooses b_{ij}):

$$m^j_{e_i}(b_{ij}) = \frac{\max_{a_i}\{m^j(a_i, b_{ij})\} + \min_{a_i}\{m^j(a_i, b_{ij})\}}{2}.$$

The perceived kindness of player *j* is given by the function

$$\lambda_{iji}(b_{ij}, c_{iji}) = m^i(b_{ij}, c_{iji}) - m^i_{e_j}(c_{iji}).$$

This is the payoff that player *i* believes that player *j* is giving to him minus the "equitable" payoff (average of maximum and minimum payoff) that player *j* can give to player *i*. Note that if player *i* expects *j* to give him less than the equitable payoff, *j*'s perceived kindness is negative, so *i* wants to give player *j* also less than the equitable payoff, and vice versa. A strategy profile $a^* = (a_i^*)_{i \in \{X, DM\}}$ is a sequential reciprocity equilibrium (SRE) if a_i^* maximizes $U^i(a_i, b_{ij}, c_{iji})$ and if $b_{ij} = a_j^*$ and $c_{iji} = a_i^*$.[30]

Proposition 3a: If *X* chooses G then DM always gets 22. If *X* chooses N then she always gets 20.

[30] See Dufwenberg and Kirchsteiger (2004) for more details and a discussion of the notion of SRE.

The equitable payoff for DM, $m_{e_X}^{DM}(b_{X,DM})$, is 21. Thus, no matter what DM believes, if *X* chooses G we have $\lambda_{DM,X,DM} = 22 - 21 = 1$, i.e., DM perceives *X*'s intentions as "kind". Similarly, if *X* chooses N we have $\lambda_{DM,X,DM} = 20 - 21 = -1$, and DM perceives *X*'s intention as "unkind."

Proposition 3b: We now show that it is an SRE if *X* chooses G and DM chooses XC.

We know already that, if producer *X* chooses G, then DM must perceive this as kind, so DM wants to reciprocate and to choose a kind action as well. By choosing action X, DM gives producer *X* a payoff of 16; by choosing C, she gives producer *X* an expected payoff of 8. The equitable payoff is (16+8)/2=12. Thus, by choosing *X*, DM gets U^{DM}(X, G, XC) = 22 + n_{DM} · (16-12) · (22-21) = 22 + 4n_{DM}; by choosing C, she obtains U^{DM}(C, G, XC) = 22 + n_{DM} · (8-12) · (22-21) = 22-4n_{DM}. Hence, for any $n_{DM} > 0$ choosing action X is optimal.

Consider now producer *X*. He believes that DM chooses the strategy XC. Furthermore, he believes that DM believes that *X* chooses G. Thus, *X* believes that DM is kind, because she reacts with *X* to G and gives him a payoff of 16 rather than 8 ($\lambda_{X,DM,X} = 16 - 12 = +4$). Therefore player *X* wants to be kind as well. If she passes on the gift U^X(G, XC, G) = 16 + n_X · (22-21) · (16-12) = 16 + 4n_X. If she does not pass on the gift she gets U^X(N, XC, G) = 9 + n_X · (20-21) · (16-12) = 9 - 4n_X. Thus, for any $n_X > 0$ choosing G is indeed optimal.

Finally, we show that it is a sequential reciprocity equilibrium that *X* chooses N and DM chooses XC. We know already that if *X* chooses N then DM must perceive this as unkind, so DM wants to reciprocate and choose an unkind action as well. By choosing action C, DM gives *X* a payoff of 1+8=9; by choosing X, she gives *X* a payoff of 1+16=17. The equitable payoff is (17+9)/2=13. Thus, by choosing C, DM gets U^{DM}(C, N, XC) = 20 + n_{DM} · (9-13) · (20-21) = 20 + 4n_{DM}; by choosing X, she obtains U^{DM}(X, N, XC) = 20 + n_{DM} · (17-13) · (20-21) = 22-4n_{DM}. Hence, for any $n_{DM} > 0$ choosing C is optimal.

Consider now producer *X*. He believes that DM chooses the strategy XC. Furthermore, he believes that DM believes that *X* chooses N. Thus, *X* believes that DM is unkind, because she reacts with C to N and gives him a payoff of 9 rather than 17 ($\lambda_{X,DM,X} = 9 - 13 = -4$). Therefore player *X* wants to be unkind as well. If she does not pass on the gift, she gets U^X(N, XC, N) = 9 + n_X · (20-21) · (9-13) = 9 + 4n_X. If she passes on the gift she gets U^X(G, XC, N) = 17 + n_X · (22-21) · (9-13) = 17-4n_X. Thus, if $n_X > 1$ choosing N is indeed optimal. ***Q.E.D.***

Proof of Proposition 4: W.o.l.g. assume that, if DM is indifferent between *X* and *Y* and also $m^C(X)$

$= m^C(Y)$, then DM chooses *X*.

(i) Suppose that producer *X* passed on the gift. If *X* is the weakly better product ($\Delta \leq 0$) DM clearly chooses *X*, both in the Gift Treatment (GT) and in the No Externality Treatment. If *X* is the worse product ($\Delta > 0$) DM chooses *X* in GT if and only if

$$22 + \alpha \cdot 0 + \bar{k} \cdot \alpha \cdot 16 + \alpha \cdot m^C(X) > 22 + \alpha \cdot 16 + \bar{k} \cdot \alpha \cdot 0 + \alpha \cdot [m^C(X) + \Delta]$$

$$\Leftrightarrow \quad \bar{k} \cdot \alpha \cdot 16 > \alpha \cdot 16 + \alpha \cdot \Delta \quad \Leftrightarrow \quad \bar{k} > 1 + \frac{\Delta}{16}$$

And in the No Externality Treatment (NET), DM chooses *X* if and only if

$$m^{DM}(X) + 2 + \alpha \cdot 0 + \bar{k} \cdot \alpha \cdot 16 > m^{DM}(X) + \Delta + 2 + \alpha \cdot 16 + \bar{k} \cdot \alpha \cdot 0$$

$$\Leftrightarrow \quad \bar{k} \cdot \alpha \cdot 16 > \alpha \cdot 16 + \Delta \quad \Leftrightarrow \quad \bar{k} > 1 + \frac{\Delta}{16\alpha}.$$

Hence, if product *X* is strictly worse ($\Delta > 0$), DM still chooses *X* for large enough $\bar{k}$, and she is more likely to do so in GT than in NET since the $\bar{k}$-threshold is lower: $1 + \frac{\Delta}{16} < 1 + \frac{\Delta}{16\alpha}$ for all $0 < \alpha < 1$.

(ii) Suppose now that producer *X* did not pass on the gift. Then, in the Gift Treatment (GT), DM chooses *X* if and only if

$$20 + \underline{k} \cdot \alpha \cdot 16 + \alpha \cdot 0 + \alpha \cdot m^C(X) \geq 20 + \underline{k} \cdot \alpha \cdot 0 + \alpha \cdot 16 + \alpha \cdot [m^C(X) + \Delta]$$

$$\Leftrightarrow \underline{k} \cdot 16\alpha \geq 16\alpha + \alpha\Delta \quad \Leftrightarrow \quad \underline{k} \geq 1 + \frac{\Delta}{16}$$

if *X* is the (weakly) better product ($\Delta \leq 0$). The inequality cannot hold if *X* is the worse product ($\Delta > 0$) because $0 \leq \underline{k} \leq 1$.

In the No Externality Treatment (NET), DM chooses *X* if and only if

$$m^{DM}(X) + \underline{k} \cdot \alpha \cdot 16 + \alpha \cdot 0 \geq m^{DM}(X) + \Delta + \underline{k} \cdot \alpha \cdot 0 + \alpha \cdot 16$$

$$\Leftrightarrow \underline{k} \cdot 16\alpha \geq \Delta + 16\alpha \quad \Leftrightarrow \quad \underline{k} \geq 1 + \frac{\Delta}{16\alpha}$$

if *X* is the (weakly) better product ($\Delta \leq 0$). Again, the inequality cannot hold if *X* is the worse product ($\Delta > 0$).

If product *X* is weakly better than *Y* ($\Delta \leq 0$), then, with $0 \leq \underline{k} \leq 1$, the inequalities hold for large enough $\underline{k}$ and, with $1 + \frac{\Delta}{16} \geq 1 + \frac{\Delta}{16\alpha}$ for $\Delta \leq 0$, it is less likely to hold in the GT than in the NET.

Q.E.D.

Proposition 5. Consider a decision maker in the Incentive Treatment (IT) who has outcome-based social preferences over an endogenously formed reference group.

(i) Suppose that producer *X* did pass on the gift. If product *X* is (weakly) better than product *Y* ($\Delta \leq 0$), DM always chooses *X* in the IT (as she does in the GT), whether or not the client shares profits. If product *X* is strictly worse ($\Delta > 0$), DM may still choose *X* in the IT, but she is less likely to do so than in the GT if the client does share profits. She is more likely to do so than in the GT if the client does not share profits.

(ii) Suppose that producer *X* did not pass on the gift. If *X* is (weakly) better than *Y* ($\Delta \leq 0$), then DM might still choose *X*, and she is more likely to do so in the IT with profit sharing than in the GT, but she is less likely to do so than in the GT if the client did not share profits. If *X* is strictly worse than product *Y* ($\Delta > 0$), instead, DM never chooses *X*, whether or not the client shares profits.

Proof of Proposition 5: W.o.l.g. assume that, if DM is indifferent between *X* and *Y* and also $m^C(X) = m^C(Y)$, then DM chooses *X*.

(i) Suppose that the client offered profit sharing and the producer *X* passed on the gift. Clearly, if *X* is at least weakly better than *Y* (i.e. $\Delta \leq 0$) then DM chooses *X* which increases her own material payoff and is good for the gift giver and the client. If $\Delta > 0$ DM chooses *X* if and only if

$$\begin{aligned} & 22 + 0.1 \cdot m^C(X) + \bar{k} \cdot \alpha \cdot 16 + \alpha \cdot 0 + \bar{k} \cdot \alpha \cdot 0.95 \cdot m^C(X) \\ & \quad > 22 + 0.1 \cdot (m^C(X) + \Delta) + \bar{k} \cdot \alpha \cdot 0 + \alpha \cdot 16 + \bar{k} \cdot \alpha \cdot 0.95 \cdot (m^C(X) + \Delta) \\ \Leftrightarrow \quad & \bar{k} \cdot \alpha \cdot 16 > 0.1 \cdot \Delta + \alpha \cdot 16 + \bar{k} \cdot \alpha \cdot 0.95 \cdot \Delta \\ \Leftrightarrow \quad & \bar{k} \cdot \alpha \cdot (16 - 0.95 \cdot \Delta) > 16\alpha + 0.1\Delta \\ \Leftrightarrow \quad & \bar{k} > \frac{16\alpha + 0.1 \cdot \Delta}{16\alpha - 0.95\alpha\Delta} = 1 + \frac{0.95\alpha + 0.1}{16\alpha - 0.95\alpha\Delta}\Delta \end{aligned}$$

If the client has not offered profit sharing and the gift was given DM still chooses *X* if *X* is at least weakly better than *Y*. He also chooses *X* if *X* if $\Delta > 0$ if and only if

$$\begin{aligned} & 22 + \bar{k} \cdot \alpha \cdot 16 + \alpha \cdot 0 + \underline{k} \cdot \alpha \cdot m^C(X) > 22 + \bar{k} \cdot \alpha \cdot 0 + \alpha \cdot 16 + \underline{k} \cdot \alpha \cdot (m^C(X) + \Delta) \\ \Leftrightarrow \quad & \bar{k} \cdot \alpha \cdot 16 > \alpha \cdot 16 + \underline{k} \cdot \alpha \cdot \Delta \\ \Leftrightarrow \quad & \bar{k} > 1 + \underline{k} \cdot \frac{\Delta}{16}. \end{aligned}$$

Hence, if the gift was given and product *X* is weakly better ($\Delta \leq 0$), then DM will always choose *X*. If product *X* is strictly worse ($\Delta > 0$), both inequalities still hold for large enough $\bar{k}$, small

enough Δ, small enough α, or small enough $\underline{k}$. To see that DM is less likely to choose *X* in the IT with profit-sharing than in the GT, we compare the thresholds for $\overline{k}$:

$$1+\frac{0.95\alpha+0.1}{16\alpha-0.95\alpha\Delta}\Delta>1+\frac{\Delta}{16}$$
$$\Leftrightarrow \frac{0.95\alpha+0.1}{16\alpha-0.95\alpha\Delta}>\frac{1}{16}$$
$$\Leftrightarrow 0.95\cdot 16\alpha+1.6>16\alpha-0.95\alpha\Delta$$
$$\Leftrightarrow 1.6>0.8\alpha-0.95\alpha\Delta$$

Since the threshold in the IT with profit-sharing is strictly larger, DM is less likely to choose *X*. Similarly, to see that she is more likely to choose *X* in the IT without profit-sharing than in the GT, we compare again thresholds and note that $1+\underline{k}\cdot\frac{\Delta}{16}\leq 1+\frac{\Delta}{16}$, with the inequality holding strictly for $\underline{k}<1$.

(ii) Suppose now that producer *X* does not pass on the gift. Then DM clearly chooses *Y* in the IT if *Y* is at least weakly better than *X*, no matter whether profit sharing has been offered or not. If *X* is strictly better than *Y* (Δ < 0) and the client offered profit sharing, DM chooses *X* if and only if

$$\begin{aligned}&20+0.1\cdot m^C(X)+\underline{k}\cdot\alpha\cdot 16+\alpha\cdot 0+\overline{k}\cdot\alpha\cdot 0.95\cdot m^C(X)\\&\quad>20+0.1\cdot(m^C(X)+\Delta)+\underline{k}\cdot\alpha\cdot 0+\alpha\cdot 16+\overline{k}\cdot\alpha\cdot 0.95\cdot(m^C(X)+\Delta)\\\Leftrightarrow\quad&\underline{k}\cdot 16\alpha>0.1\Delta+16\alpha+\overline{k}\cdot\alpha\cdot 0.95\Delta\\\Leftrightarrow\quad&\underline{k}>1+\frac{0.95\alpha\overline{k}+0.1}{16\alpha}\Delta.\end{aligned}$$

If Δ < 0 and the client has not offered profit sharing, DM chooses *X* if and only if

$$\begin{aligned}&20+\underline{k}\cdot\alpha\cdot 16+\alpha\cdot 0+\underline{k}\cdot\alpha\cdot m^C(X)>20+\underline{k}\cdot\alpha\cdot 0+\alpha\cdot 16+\underline{k}\cdot\alpha\cdot(m^C(X)+\Delta)\\\Leftrightarrow\quad&\underline{k}\cdot 16\alpha>16\alpha+\underline{k}\cdot\alpha\cdot\Delta\\\Leftrightarrow\quad&\underline{k}\cdot(16-\Delta)>16\\\Leftrightarrow\quad&\underline{k}>\frac{16}{16-\Delta}.\end{aligned}$$

Hence, if the gift was not given and product *X* is strictly worse (Δ > 0), then DM will never choose *X*. If instead product *X* is strictly better (Δ < 0), both inequalities hold for large enough $\underline{k}$, large enough Δ, large enough α, or small enough $\overline{k}$, and DM may choose *X*. To see that DM is more likely to choose *X* in the IT with profit-sharing than in the GT, we compare the thresholds for $\underline{k}$:

$$1+\frac{0.95\alpha\bar{k}+0.1}{16\alpha}\Delta<1+\frac{1}{16}\Delta$$
$$\Leftrightarrow \frac{0.95\alpha\bar{k}+0.1}{\alpha}>1$$
$$\Leftrightarrow 0.95\alpha\bar{k}+0.1>\alpha.$$

Since the threshold in the IT with profit-sharing is strictly smaller, DM is more likely to choose *X*. Similarly, she is less likely to do so in the IT without profit-sharing than in the GT since, for $\Delta < 0$,

$$\frac{16}{16-\Delta}>\frac{16-\Delta}{16}$$
$$\Leftrightarrow 16^2>(16-\Delta)^2.$$

Q.E.D.

Appendix B: Experimental Parameterization

Period	Potential gift giver	Possible payoffs of product *A*		Expected value of *A*	Spread btw. payoffs of *A*	Possible payoffs of product *B*		Expected value of *B*	Spread btw. payoffs of *B*	Diff. in EVs (pot. gift giver minus other)	Diff. in Spreads (pot. gift giver minus other)
		50%	50%			50%	50%				
1	*A*	13	15	14	2	20	12	16	8	-2	-6
2	*B*	15	17	16	2	12	20	16	8	0	6
3	*B*	16	14	15	2	14	20	17	6	2	2
4	*B*	13	19	16	6	5	15	10	10	-6	4
5	*A*	17	7	12	10	10	14	12	4	0	6
6	*B*	12	16	14	4	19	13	16	6	2	-6
7	*A*	11	19	15	8	18	16	17	2	-2	2
8	*A*	8	20	14	12	10	18	14	8	0	4
9	*B*	17	19	18	2	10	14	12	4	-6	2
10	*A*	19	13	16	6	20	8	14	12	2	-6
11	*B*	20	12	16	8	7	13	10	6	-6	-2
12	*B*	3	17	10	14	5	11	8	6	-2	-8
13	*A*	16	12	14	4	8	20	14	12	0	-8
14	*A*	9	15	12	6	19	5	12	14	0	-8
15	*B*	19	11	15	8	7	19	13	12	-2	4
16	*A*	8	12	10	4	13	3	8	10	2	-6
17	*B*	20	16	18	4	16	8	12	8	-6	4
18	*A*	7	13	10	6	16	8	12	8	-2	-2
19	*A*	8	14	11	6	14	12	13	2	-2	4
20	*B*	13	19	16	6	18	14	16	4	0	-2
min		3	7	10	2	5	3	8	2	-6	-8
max		20	20	18	14	20	20	17	14	2	6
avg		13.20	15.00	14.10	6.00	13.05	13.15	13.10	7.50	-1.40	-0.8

TABLE A1: Payoffs of the different products in the 20 periods

There are

- four periods in which the potential gift giver's expected value is 2 points higher (periods 3, 6, 10, and 16)
- six periods in which there is no difference in expected value between producer *A* and producer *B* (periods 2, 5, 8, 13, 14, and 20)
- six periods in which the potential gift giver's expected value is 2 points lower (periods 1, 7, 12, 15, 18, 19)
- four periods in which the potential gift giver's expected value is 6 points lower (periods 4, 9, 11, 17)

Note that in the four periods in which the potential gift giver's expected value is 6 points lower, his lottery is first order stochastically dominated by the lottery of his competitor.

Note further that there are 10 periods in which the spread between possible payoffs is higher for the product of the potential gift giver than for the alternative product, and 10 periods in which it is lower. Among the six periods with equal expected values, the spread is larger in three periods and lower in the other three periods.

Table 1. Summary Statistics by Treatment

	Observations					Demographics							
	Total	Subjects	DMs	Producers	Clients	Gender:	Age			Major			
						Females	20s	30s	40s-60s	Econ/Bus	Other Soc Sc	Human.	Nat. Sc.
1. **Baseline Treatment (BT)**	240	24	12	0	12	46%	92%	8%	0%	42%	4%	29%	17%
Producers cannot make any gifts.													
2. **Gift Treatment (GT)**	480	96	24	48	24	55%	96%	3%	1%	26%	11%	23%	26%
One producer can give a gift.													
3. **No Externality Treatment (NET)**	320	48	16	32	0	46%	96%	4%	0%	23%	19%	17%	23%
One producer can give a gift.													
No client; DM is residual claimant.													
4. **Incentive Treatment (ICT)**	460	92	23	46	23	55%	96%	4%	0%	28%	11%	15%	16%
One producer can give gift.													
Client can offer profit sharing.													
5. **Large-Gift Treatment (LGT)**	240	48	12	24	12	63%	98%	2%	0%	19%	13%	10%	25%
One producer can give large gift.													
6. **Disclosure Treatment (DT)**	240	48	12	24	12	54%	94%	4%	2%	13%	17%	23%	21%
One producer can give gift.													
Client informed about gift and DM's response.													
Total	1,980	356	99	174	83	54%	96%	4%	1%	24%	13%	19%	22%

Table 2. Gift Treatment

	OLS	OLS	Logit	OLS	OLS	Logit
	(1)	(2)	(3)	(4)	(5)	(6)
Gift given	0.338***	0.334***	0.336***			
	(0.060)	(0.061)	(0.062)			
Gift not given	-0.130**	-0.131**	-0.164**			
	(0.057)	(0.057)	(0.074)			
Product *X* has higher EV [+2]				0.500***	0.503***	0.546***
				(0.053)	(0.053)	(0.057)
Product *X* has lower EV [-2]				-0.361***	-0.364***	-0.533***
				(0.039)	(0.041)	(0.069)
Product *X* has lower EV [-6]				-0.333***	-0.333***	-0.429***
				(0.061)	(0.061)	(0.100)
(Product *X* has same EV)*(Gift given)				0.436***	0.433***	0.440***
				(0.062)	(0.065)	(0.056)
(Product *X* has same EV)*(Gift not given)				-0.155*	-0.158*	-0.174**
				(0.082)	(0.079)	(0.086)
(Product *X* has higher EV [+2])*(Gift given)				0.011	0.010	0.040
				(0.068)	(0.069)	(0.244)
(Product *X* has higher EV [+2])*(Gift not given)				-0.361***	-0.369***	-0.409***
				(0.120)	(0.126)	(0.082)
(Product *X* has lower EV [-2])*(Gift given)				0.415***	0.413***	0.516***
				(0.091)	(0.091)	(0.088)
(Product *X* has lower EV [-2])*(Gift not given)				0.016	0.011	0.056
				(0.050)	(0.048)	(0.200)
(Product *X* has lower EV [-6])*(Gift given)				0.402***	0.393***	0.452***
				(0.119)	(0.117)	(0.118)
(Product *X* has lower EV [-6])*(Gift not given)				-0.083	-0.070	
				(0.063)	(0.059)	
Female		0.050	0.062		0.0486	0.0773
		(0.044)	(0.055)		(0.045)	(0.075)
Economist		0.0612	0.0740		0.0608	0.1000
		(0.045)	(0.055)		(0.046)	(0.076)
Period		-0.00497*	-0.00595*		0.00174	0.00285
		(0.003)	(0.003)		(0.003)	(0.004)
Constant	0.342***	0.351***		0.417***	0.356***	
	(0.032)	(0.059)		(0.047)	(0.069)	
Observations	720	720	720	720	720	694
(Pseudo) R-squared	0.156	0.164	0.124	0.407	0.412	0.331

Notes. The sample consists of the Baseline and the Gift Treatments. *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. The table reports coefficients of OLS regressions in columns (1), (2), (4), and (5), and marginal effects of logit regressions in columns (3) and (6). Standard errors are adjusted for clustering at the decision-maker level. The interaction term (Product *X* has lower EV [-6])*(Gift not given) perfectly predicts the outcome in logit regression (6).

*** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

Table 3. 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	0.218***	0.210***	-0.125*			
	(0.049)	(0.052)	(0.064)			
NET: Gift not given	-0.036	-0.043	0.087			
	(0.047)	(0.050)	(0.059)			
(Product *X* has same EV)*(NET: Gift given)				0.357***	0.350***	-0.084
				(0.067)	(0.070)	(0.064)
(Product *X* has same EV)*(NET: Gift not given)				-0.045	-0.053	0.104
				(0.093)	(0.093)	(0.101)
(Product *X* has higher EV [+2])*(NET: Gift given)				0.099**	0.091**	0.080
				(0.038)	(0.039)	(0.061)
(Product *X* has higher EV [+2])*(NET: Gift not given)				-0.080	-0.086	0.280**
				(0.066)	(0.069)	(0.130)
(Product *X* has lower EV [-2])*(NET: Gift given)				0.258***	0.248***	-0.166
				(0.086)	(0.085)	(0.114)
(Product *X* has lower EV [-2])*(NET: Gift not given)				0.039	0.035	0.023
				(0.050)	(0.055)	(0.059)
(Product *X* has lower EV [-6])*(NET: Gift given)				-0.032	-0.038	-0.433***
				(0.067)	(0.072)	(0.106)
(Product *X* has lower EV [-6])*(NET: Gift not given)				-0.068	-0.076	-0.005
				(0.054)	(0.058)	(0.012)
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, GT, BT	NET, BT	NET, GT, BT	NET, GT, BT
Observations	560	1,040	1,040	560	1,040	1,040
R-square	0.055	0.140	0.140	0.492	0.440	0.440

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 abbreviations *gg* and *gng* indicate *Gift given* and *Gift not given*, respectively. Constant included. Standard errors are clustered at the decision maker level. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

Table 4. Policy Treatments: Disclosure, Size, and Financial Incentives

	Policy: Disclosure		**Policy: Large Gift**		**Policy: Incentive-ps**		**Policy: Incentive-nps**	
	Diff. to BT	*Diff. to GT*	*Diff. to BT*	*Diff. to GT*	*Diff. to BT*	*Diff. to GT*	*Diff. to BT*	*Diff. to GT*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Model 1: Overall Effects								
Policy Treatment: Gift given	0.318***	-0.014	0.163**	-0.165*	0.263***	-0.071	0.414***	0.075
	(0.073)	(0.082)	(0.076)	(0.084)	(0.067)	(0.076)	(0.065)	(0.077)
Policy Treatment: Gift not given	-0.120*	0.014	-0.077	0.065	-0.033	0.097	-0.097*	0.026
	(0.065)	(0.069)	(0.072)	(0.072)	(0.088)	(0.094)	(0.052)	(0.062)
Dummies for (GT: Gift given) and (GT: Gift not given)	X		X		X		X	
Dummies for (Gift given) and (Gift not given)		X		X		X		X
Controls for gender, major, and period	X	X	X	X	X	X	X	X
R-square	0.166	0.166	0.128	0.128	0.153	0.153	0.195	0.195
Model 2: Estimates by EV differences								
(Product *X* has same EV)*(Policy T: Gift given)	0.375***	-0.056	0.304***	-0.124	0.414***	-0.021	0.398***	-0.041
	(0.090)	(0.084)	(0.100)	(0.093)	(0.078)	(0.072)	(0.080)	(0.077)
(Product *X* has same EV)*(Policy T: Gift not given)	-0.076	0.087	-0.208	-0.039	-0.178	-0.021	-0.163	-0.015
	(0.205)	(0.206)	(0.138)	(0.140)	(0.140)	(0.146)	(0.110)	(0.119)
(Product *X* has higher EV [+2])*(Policy T: Gift given)	0.093**	0.086	-0.083	-0.089	0.032	0.022	0.046	0.032
	(0.044)	(0.065)	(0.075)	(0.087)	(0.057)	(0.074)	(0.047)	(0.066)
(Product *X* has higher EV [+2])*(Policy T: Gift not given)	-0.330*	0.044	-0.205	0.179	-0.038	0.328*	-0.759***	-0.405***
	(0.165)	(0.196)	(0.232)	(0.251)	(0.127)	(0.169)	(0.082)	(0.139)
(Product *X* has lower EV [-2])*(Policy T: Gift given)	0.443***	0.033	0.233**	-0.174	0.272**	-0.142	0.562***	0.144
	(0.107)	(0.131)	(0.114)	(0.136)	(0.114)	(0.138)	(0.097)	(0.126)
(Product *X* has lower EV [-2])*(Policy T: Gift not given)	-0.044	-0.050	0.220	0.221	0.151	0.138	0.139	0.119
	(0.036)	(0.040)	(0.290)	(0.289)	(0.114)	(0.117)	(0.098)	(0.101)
(Product *X* has lower EV [-6])*(Policy T: Gift given)	0.186	-0.204	0.116	-0.271*	0.033	-0.362***	0.592***	0.191
	(0.133)	(0.156)	(0.120)	(0.145)	(0.103)	(0.130)	(0.106)	(0.134)
(Product *X* has lower EV [-6])*(Policyt T: Gift not given)	-0.010	0.061	-0.144*	-0.065	-0.070	0.002	0.297***	0.364***
	(0.078)	(0.053)	(0.080)	(0.049)	(0.058)	(0.013)	(0.108)	(0.091)
Dummies for EV differences	X	X	X	X	X	X	X	X
Dummies for EV differences interacted with (GT: gg) and (GT: gng)	X		X		X		X	
Dummies for EV differences interacted with (gg) and (gng)		X		X		X		X
Controls for gender, major, and period	X	X	X	X	X	X	X	X
R-square	0.408	0.408	0.383	0.383	0.421	0.421	0.380	0.380
Sample	DT, GT, BT	DT, GT, BT	LGT, GT, BT	LGT, GT, BT	IT, GT, BT	IT, GT, BT	IT, GT, BT	IT, GT, BT
Observations	960	960	960	960	851	851	1,049	1,049

Notes. The DT sample contains all observations from the Disclosure Treatment; LGT the observations from the Large Gift Treatment, IT the observations from the Incentive Treatment; BT the observations from the Baseline Treatment; and GT the observations from the Gift Treatment. The abbreviations gg and gng indicate "gift given" and "gift not given," respectively; the abbreviations ps and nps indicate "profit shared" (by the client) and "no profit shared" (by the client) respectively. Constant included. Standard errors are clustered at the decision maker level. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.