You Owe Me*

Ulrike Malmendier[a)] and Klaus M. Schmidt[b)]

July 6, 2014

Abstract

In business and politics, gifts are often aimed at influencing the recipient at the expense of others. In an experimental study, which removes informational and incentive confounds, subjects strongly respond to small gifts even though they understand the gift giver's intention. Our findings question existing models of social preferences. They point to anthropological and sociological theories about gifts creating an obligation to reciprocate. We capture these effects in a simple extension of existing models. We also show that common policy responses (disclosure, size limits) are ineffective, consistent with our model. Financial incentives are effective but backfire when not used.

Keywords: Gift exchange; externalities; lobbyism; corruption; reciprocity; social preferences.

JEL: C91, D62, D73, I11.

*) We would like to thank Yan Chen, Rachel Croson, Robert Dur, Ernst Fehr, Rudolf Kerschbamer, Matthias Kräkel, George Loewenstein, Cristina Nistor. We would also like to thank seminar participants in the 2011 ASSA meetings, the Third Behavioral Economics Annual Meeting (BEAM), and at the universities of Aix-Marseilles, Berkeley, Emory, Innsbruck, Maastricht, MIT, Munich, Pennsylvania, Princeton, and Toulouse for comments and suggestions. David Arnold, Leonie Giessing, Jana Jarecki, Sara Kwasnick, Eddie Ning, Warren Ronsiek, and Slobodan Sudaric provided excellent research assistance. Financial support by Deutsche Forschungsgemeinschaft through SFB-TR 15, the Excellence Initiative of the German government, and the Alfred P. Sloan Foundation is gratefully acknowledged.
[a)] Ulrike Malmendier, Department of Economics, University of California at Berkeley, 501 Evans Hall, Berkeley, CA 94720-3880, USA, email : ulrike@econ.berkeley.edu
[b)] Klaus M. Schmidt, Department of Economics, University of Munich, Ludwigstrasse 28, D-80539 Munich, Germany, email: klaus.schmidt@LMU.de

1 Introduction

Social preferences influence behavior in many economically important settings, including bargaining, incentive provision at the workplace, and cooperation and competition in small groups (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 (Fehr, Goette and Zehnder, 2009). A well-developed body of social-preference theories provides explanations for the observed gift-exchange behavior, such as intention-based reciprocity (Rabin, 1993) or inequality 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 the behavior of the recipient, 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 a gift (ranging from small tokens of appreciation, such as pens or coffee mugs, to bottles of wine or event tickets) from a supplier, who hopes to get favorable treatment, 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 third-party "client" who is often anonymous: shareholders, the electorate, 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 with negative externalities are effective. The contribution of our paper is threefold. First, we

[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.

show that there is a powerful effect of the gift *per se*. Subjects reciprocate to gifts even if there are no monetary incentives for doing so and the gift does not convey positive information about the gift giver or about the product offered by the gift giver.[3] We observe a large reciprocal response even though the gift is small and the recipients fully understand that the gift is given for selfish reasons, namely to influence their behavior at the expense of a third party. The effect goes in both directions: the decision maker favors the potential gift giver if the gift is given and she discriminates against him if the gift is not given. Second, we show that most standard models of social preferences have difficulties predicting the observed behavior. This behavior is better predicted by anthropological and sociological theories positing that gifts create a bond between the gift giver and the recipient and an obligation to reciprocate. We propose a simple extension of standard models of social preferences that formalizes this idea. Third, we conduct three policy treatments to evaluate the effectiveness of policy interventions that are often proposed to mitigate the effects of gifts with negative externalities.

In our main treatment (*Gift Treatment*), a decision maker has to buy one of two possible products on behalf of a client. Before making the decision, she may receive a small gift from one of the two producers. The gift is given unconditionally and before the producers learn about the value of their products, so that the gift cannot contain any information about the product's quality. The setting is anonymous, and players are re-matched after each round. Neither client nor producers find out which product the decision maker picks, which excludes any reputation effects. Nevertheless the effect of the gift is large. It is significantly stronger than in the classic gift-exchange game without a third party, which we conducted in a control treatment (*No Externality Treatment*). Furthermore, we observe that the gift is given with the unambiguous intention to affect the decision of the recipient at the expense of a third party. We show that the recipients are fully aware of this intention, but they reciprocate nevertheless.

We also find that *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 that product is better. Thus, the punishment also comes at the expense of the client.

Our experimental design allows us to test whether decision makers are aware of how

[3] Of course, many gifts provide financial incentives or convey information. 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. These effects are important in many circumstances and are explained by standard economic theory. In this paper, we focus on the effects of a gift *per se*, i.e. in the absence of any incentive or information effects. Our results imply that regulation removing incentives may not suffice to eliminate over-prescription of drugs.

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. Furthermore, they somewhat overestimate by how much the other decision makers are affected by the gift, but they estimate the difference between their own and others' behavior to be much smaller than in Steinman et al. (2001).

The most prominent economic theories of other-regarding behavior have difficulties predicting the observed phenomena. Theories of outcome-based social preferences such as altruism (Andreoni and Miller, 2002), maximin preferences (Charness and Rabin, 2002) or inequality aversion (Fehr and Schmidt, 1999) and of type-based reciprocity (Levin, 1998) are refuted by the observed behavior. Theories of intention-based reciprocity (Rabin 1993, Dufwenberg and Kirchsteiger, 2004) are consistent with our experimental results. They admit an equilibrium in which the gift is given and the decision maker chooses the gift giver. However, they also admit an equilibrium in which no gifts are given and the decision maker always chooses the best product for the client. Thus, they have little predictive power. Furthermore, the underlying motives seem different. These theories are based on the idea that the gift affects social preferences by signaling that the gift giver has kind intentions. In our experiment, however, producers admit in a questionnaire that they give the gift for purely selfish reasons, and the decision makers state that they perfectly understand this.

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

We propose a simple extension of existing models of social preferences that formalizes this idea. In the model the weight that player *i* attaches to the welfare of player *j* depends on the

actions of j that affect i, relative to the expected behavior of j. A favorable act such as giving a gift strengthens the bond between the giver and the recipient, i.e., the weight of the giver's payoff in the recipient's utility function, and the recipient will reciprocate. The more the favorable act exceeds expectation, the stronger the positive response. This simple model has a unique equilibrium that 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 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. When we inform clients which producer is the potential gift giver, whether he chooses to give the gift, and which product the decision maker chooses, decision makers' behavior remains very similar even though they know that everything is disclosed.[4]

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 53 percent of all cases, compared to 68 percent when the gift is small. This is contrary to the logic of size limits, and may be surprising at first glance, but it is predicted by our model: decision makers correctly expect that producers will send the gift with a higher probability if it benefits them more. The more the gift is expected, the smaller is the effect on behavior.

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

The rest of the paper is organized as follows: The next section discusses the relation of our paper to the literature. Section 3 describes the experimental design. 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 Treat-

[4] Note, however, that our design does not allow clients to react to the decision maker's behavior, e.g., by choosing another decision maker, because we restrict attention to one-shot interaction. A design that allows for repeated interaction may lead to different results.

ments. 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 6), our paper is related to three branches of the economics literature. First, there is a large experimental literature on gift exchange games, starting with Fehr, Kirchsteiger and Riedl (1993). This literature has established reciprocity as an important motive facilitating gift exchange. In almost all of this literature, however, the gift affects only the giver and the receiver; there are no externalities. A notable exception 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) penalty and staff rotation significantly reduces corruption. In contrast, 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.[5]

A second related strand of the literature studies the effects of gifts with positive or negative externalities in field experiments. Currie, Lin and Meng (2012) conducted a field experiment in Chinese hospital outpatient clinics. In the experiment, two trained auditors acting as patients visit a physician in sequence. If the first patient gives a small gift to the physician, he receives better service and he is less likely to be prescribed unnecessary and costly antibiotics than if no gift is given. Furthermore, if the first patient introduces the second patient as his friend this patient also receives better service. If the second patient is not introduced, he is treated significantly worse. 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)

[5] Our evidence on gifts triggering an obligation to "repay" independently of giver's intentions is also consistent with behavior observed in a modified trust game, in which 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, Strassmair shows that the behavior of trustees is unaffected by the perceived kindness of the trustor. As in our results, intentions do not seem to play a role.

estimate the impact of a large anti-poverty program in Uruguay on political support for the government that implemented it. Those (quasi-)randomly selected households that benefited from the program are 11 to 13 percentage points more likely to favor the current government than those who did not benefit. These studies show that reciprocity effects with externalities are not restricted to the lab but extend to the field. However, these studies cannot identify the behavioral mechanism that is driving the observed behavior.

Third, a large empirical literature studies business gifts, especially in the pharmaceutical industry. In a meta-study of 29 data sets, Wazana (2000) concludes that gifts are "associated with increased prescription rates of the sponsor's medication" (p. 373). Campbell et al. (2007) conduct 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 among physicians on whether it is ethical to accept gifts of the pharmaceutical industry and whether these gifts affect prescriptions. The general conclusion in this literature is that business gifts are widespread and effective. However, as discussed in Dana and Loewenstein (2003), the empirical literature cannot disentangle the causal factors that explain why gifts work.

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

In the experiment, two producers want to sell their products to a client. At the beginning of each period one producer is chosen randomly as the potential gift giver while the other producer cannot take any action. The potential gift giver, called producer X, offers product X; the other producer, called producer Y, offers product Y. The client has to buy either product X or product Y but has to rely on an expert to make this decision on his behalf. We call the expert the decision maker (DM). DM receives a fixed wage for her services and is instructed to choose the best product for the client. If the decision maker chooses product X (product Y, respectively), producer X (Y) receives a positive (quasi-)rent, while the other producer gets 0. Before DM takes her decision producer X can pass a small monetary gift to her. The gift is unconditional, there is

no possibility to refuse the gift, and there is no repeated interaction – subjects are anonymously re-matched after each round. The client is aware of the possibility of a gift, but does not know whether a gift was actually given, nor which product was chosen.

We implemented this set-up as follows. A typical session has 24 subjects: 6 decision makers, 6 clients, 6 producers *X*, and 6 producers *Y*.[6] There are 20 periods. In each period the decision maker is anonymously matched with a new client and new producers *X* and *Y*. Producer *X* receives one token from the experimenter that he can keep or pass on to the decision maker in which case it doubles and DM receives two tokens. Products *X* and *Y* are simple 50/50 lotteries. The payoffs are natural numbers between 3 and 20. For example, product *X* may pay out 5 or 11 with equal probability, while product *Y* pays 3 or 17. The lottery pairs in each period are fixed throughout all treatments and sessions. In four periods, the expected value 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 *X* is 2 points lower than the expected value of *Y*. In four periods, the expected value of *X* is 6 points lower. In this last case, product *Y* first-order stochastically dominates product *X*, i.e., every rational decision maker (no matter how risk averse or risk loving) prefers *Y* to *X*. (Appendix-Table A1 in Appendix B shows all 20 lottery pairs.)

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 the decision

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

maker actually chooses, 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.

In each session the instructions are read aloud to make sure that the rules of the experiment and the information structure are common knowledge between all subjects.

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 the 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 tries 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) described in Section 7.

Discussion of Design Features

The experimental design has several features that merit discussion. First, only one of the two producers can make a gift. Had we allowed both producers to offer gifts, both producers would likely have done so in most periods, in which case we would have lost many observations. Allowing both producers to offer a gift would have been more interesting in a set up where producers can also choose the size of the gift. This would have turned the game into a gift giving contest, giving rise to many additional interesting questions that we have to leave to future research.

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

Third, the products are simple lotteries over monetary outcomes. We chose this design feature to give the decision maker some moral wiggle room: Even when product *Y* has a higher expected value than product *X*, the decision maker may justify buying *X* because it is less risky or because it gives a higher payoff in one state of the world than *Y*. At the same time, our design does include cases in which every rational agent prefers product *Y*: If the difference in expected values is 6, then *Y* first order stochastically dominates *X*. We also note that such probabilistic consequences are an important feature of many real-world decisions on behalf of a third party. For example, the effects of a drug or a policy measure are typically stochastic.

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

Experimental procedures

Overall, we conducted 15 sessions with 20-24 participants in each session at MELESSA[7] 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. 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. 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 €13.20 ($18.04), which includes a show-up fee of €4 ($5.47). Further summary statistics are in Appendix-Table A2.

4 Behavioral Predictions

To guide our empirical analysis, we present a simple theoretical framework that allows us to de-

[7] MELESSA is the Munich Experimental Laboratory for the Economic and Social Sciences. We used the software ORSEE (Greiner 2004) for recruitment and z-Tree (Fischbacher 2007) for the experiments.

rive the predictions of existing theories of social preferences and reciprocity.

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 gift is given unconditionally and prior to her decision. Thus, the traditional model of rational and self-interested behavior does not predict that the decision maker favors the gift giver. She is indifferent, no matter whether or not the gift was given. If there is no client and the decision maker chooses the product for herself, the traditional model predicts that she chooses the product with 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.[8] To sum up the main insights of this section, we show that, in the context of our experiment, outcome-based theories (e.g. inequality aversion or altruism) predict unambiguously that decision makers should *not* be influenced by gifts but should maximize the expected utility of their clients. This is contradicted by our experimental results. Models of type-based reciprocity are not consistent with our data either. In any pooling equilibrium the decision maker should favor the client and in any (partially) separating equilibrium passing on the gift is a signal of kindness only if the reaction of DMs to the gift is much smaller than what we observe. Models of intention-based reciprocity have multiple equilibria. There is an equilibrium in which all producers pass on the gift and the decision makers favor the gift giver. However, there is also an equilibrium in which no gifts are given and the decision makers favor their clients. Thus, intention-based reciprocity is consistent with the main observations of the experiment, but it lacks predictive power. The rest of this section makes these claims precise.

We assume that the decision maker is risk neutral and evaluates products *X* and *Y* by their expected values.[9] 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 in-

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

[9] 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 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.

different she chooses the product that favors the client.[10]

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 $i \in \{DM, X, Y, C\}$ 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 (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 (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.[11]

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 producer 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 pay-

[10] 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.

[11] 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.

off 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 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*.[12]

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 they 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 ex-ante probability of being kind is μ, i.e., $\mu^{DM} = \mu^X = \mu^Y = \mu^C = \mu$, 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 *i* is given by

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

where $\alpha > 0$ is the (common) degree to which a kind player *i* cares about the payoff of a kind player j.[13] The utility function of a selfish player *i* simply is $U^{i,selfish} = m^i$.

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

[12] 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 then she favors *Y*; else she favors the client.

[13] We could also assume that a kind player cares about the payoff of a selfish player to the degree $\underline{\alpha}$ and about the payoff of a kind player to the degree $\overline{\alpha}$, 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.

Proof: See Appendix A.

In our experiments the decision makers favor the producer *X* if the gift was given and producer *Y* if the gift was not given. This is inconsistent with a pooling equilibrium. It is also inconsistent with a (partially) separating equilibrium, because the selfish type of the producer has a strong incentive to mimic the kind type. In our experiment the probability that DM chooses *X* increases by 47 percent if the gift is given, much more than 1/16=6.25 percent. Thus, producer *X* cannot signal that he is a kind type by giving the gift.

Finally, we consider 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. 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 (SRE)" of Dufwenberg and Kirchsteiger (2004) that generalizes Rabin's (1993) notion of "Fairness Equilibrium" to sequential games.

> **Proposition 3** Suppose that DM and producer *X* are motivated by intention-based reciprocity. If they care strongly enough about the kindness of the intentions of the other player, then 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: See Appendix A.

Psychological games with intention-based reciprocity are consistent with many interesting phenomena, but they are also plagued by multiple equilibria. It is an equilibrium that both players are kind to each other if both expect the other player to be kind as well. But it is also an equilibrium that both players behave unkindly if they believe that the other player is hostile.[14]

[14] 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.

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 when 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 20 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.

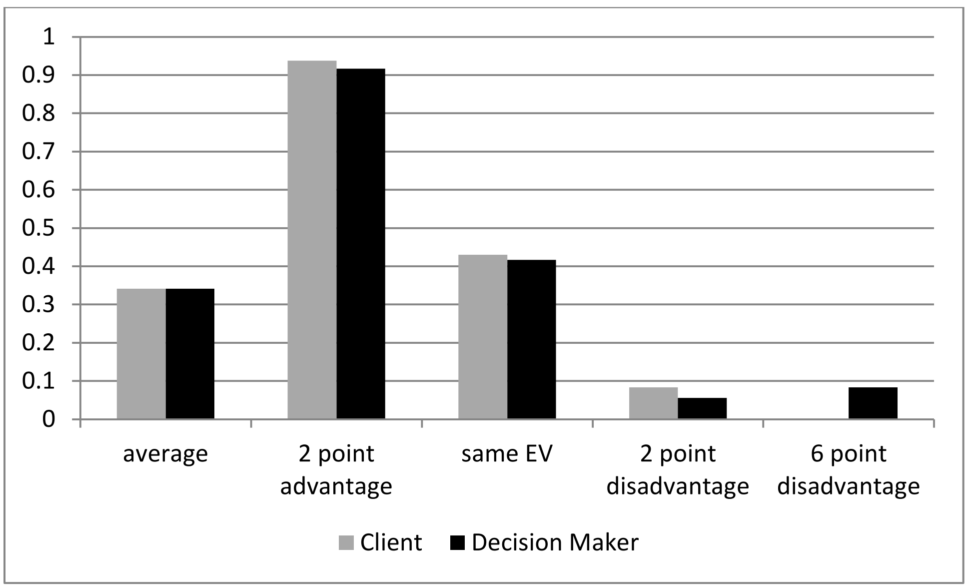

FIGURE 1: Frequencies of choosing *X* in the Baseline Treatment

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. If 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.[15] 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 (or the same expected value and the lower variance) in only 7 out of 20 periods, i.e. in 35 percent of all cases. Also note that a little less than 50 percent of clients and 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 much weaker preference for the product with the lower variance. This is why we ignore the (insignificant) effects of the variance in the following.

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 that the average effect of the gift is large: 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 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.[16]

[15] 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 ($p = 0.906$ and $p = 0.867$).

[16] Fligner-Policello robust rank-order tests comparing the average decisions of each DM in the Gift Treatment when the gift was given (not given) to the average decisions of each DM in the Baseline Treatment show that these differences are significant with $p < .001$ ($p = .023$). We use the Fligner-Policello robust rank order-test rather than the Wilcoxon-Mann-Whitney test to account for unequal variances across populations. This is important in our analysis as the number of periods we average over varies. For example, in the Baseline Treatment, we calculate the frequency of choosing *X* over 20 periods. But in the Gift Treatment we calculate separately the average decisions when the gift was given and when it was not given and hence use fewer periods; the variance is mechanistically larger.

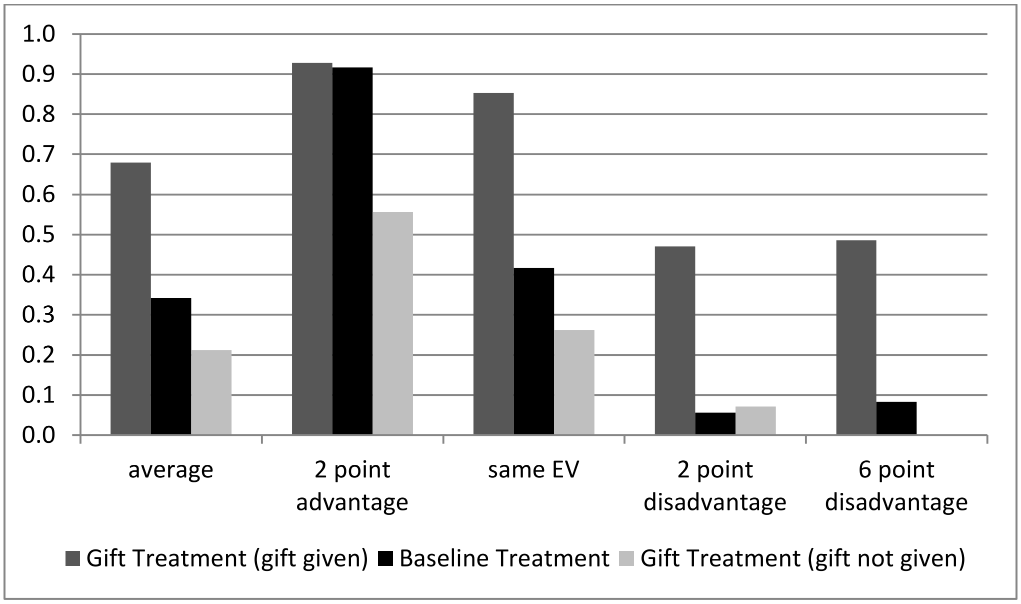

FIGURE 2: Frequencies of choosing *X* in the Gift Treatment (with comparison to the Baseline Treatment)

The next set of bars shows that, if the lottery 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 even though it is the better product. Over 35 percent of DMs "punish" producers for not giving the gift, at the expense of the client.

If the expected value of the two products is the same, we observe both the positive and the negative effect of the gift: the fraction choosing *X* doubles, compared to the Baseline, if the gift was passed on, from 41.7 percent to 85.3 percent; but it 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 is the preferred product of the client. Nevertheless, *X* is chosen almost half of the time when the gift was given.[17]

[17] Fligner-Policello robust rank-order tests comparing the **average** decisions of DMs in the Gift Treatment when the gift was given to the average decisions of DMs in the Baseline Treatment show that these decisions are significantly different if product *X* has the same expected value ($p <. 001$), a two point disadvantage ($p <. 001$) or a six point dis-

In Table 1, 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 both a linear probability model (OLS) and maximum-likelihood estimation (logit). The regression framework allows us to control for demographics and period-effects and to account for heteroskedasticity, within-subject correlation, and within-session correlations using robust, cluster-robust, wild cluster bootstrap-t procedures. In the main table, we report robust standard errors. In addition, we calculate cluster-robust standard errors accounting for within-subject correlation, and we also re-estimated the coefficients using fixed effects, both with robust standard errors and with clustered standard errors. The regression results are almost identical. Finally, to account for within-session correlations, we cannot cluster by session as the number of clusters is too small. We implement the *wild cluster bootstrap-t* procedure of Cameron, Gelbach, and Miller (2008) to deal with cases of few clusters (see Appendix C for details). We computed the significance levels using 10,000 repetitions. Appendix-Table A3 displays the bootstrap results for session-level clustering. For completeness (and also because the number of DMs is sometimes too small to apply standard asymptotic tests), we also wild cluster bootstrap assuming DM clusters. (The latter results are very similar to cluster-robust standard errors.)

Columns 1 to 3 of Table 1 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 not giving a gift (when gift giving is possible), reduces the probability of *X* being chosen by about one third. Both effects are highly significant ($p < 0.01$).[18] Adding controls does not affect the coefficient estimates. In particular, neither gender nor the field of study (economics/business versus others) has a significant effect. Interacting gender and major, respectively, with receiving and not receiving gifts is similarly inconsequential, with one exception: The estimate on the interaction between female and not receiving a gift is positive ($p <.05$). The *Period* coefficient indicates a slightly negative time trend in decision makers' 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,

---

advantage ($p = .034$). When the gift was not given in the Gift Treatment, the difference is statistically significant if *X* has a two point disadvantage ($p = .050$) or the same expected value ($p =. 095$).

[18] Cluster-robust standard errors and wild cluster bootstrapping for individual-level clusters yield similar results. When bootstrapping with session-level clusters, the effects of (not) giving is significant with $p = .031$ ($p = .093$) in the baseline specification (Appendix-Table A3, col. 1). The effect of giving remains significant at the same level when adding controls (col. 2), but the effect of not giving becomes insignificant ($p = .153$).

this slight trend comes from the Baseline Treatment, where no gift giving is possible; conditional on receiving or not receiving the gift, there are no time trends in the Gift Treatment.[19]

In the next three columns, we refine the regression model to account for the differences in expected values between the two products, as we did in the second to fifth set of bars in Figure 2. The intercept captures the case of "no difference in expected value" in the 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 …") go in the expected direction: 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). Each effect is significant at the one-percent level ($p < 0.01$). Hence, in a world without gifts, the influence of product value is economically and statistically highly significant

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

Overall, the Gift Treatment shows that gifts can have strong externalities. Clients receive the worse product with almost 50 percent probability (instead of less than 10 pct) if the producer of the worse product has sent a gift or 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 to a setting without externalities. If gifts have the same effect, regardless of who pays the cost, then there is no obvious distortion and it is more difficult to argue that gifts induce 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

[19] 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. 0.005) 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).

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 send the gift.

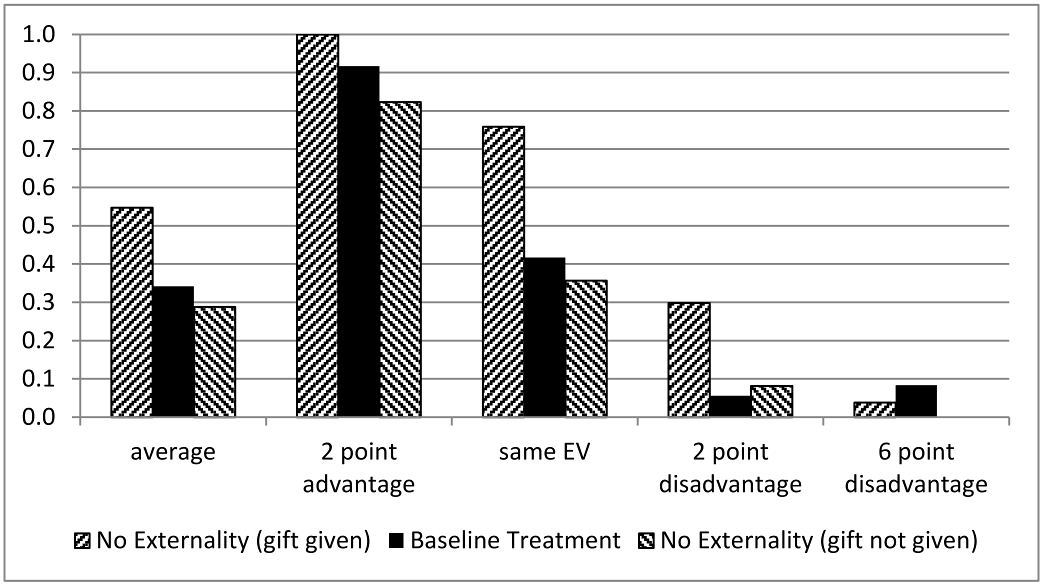

FIGURE 3: Frequencies of choosing *X* in the No Externality Treatment (with comparison to the Baseline Treatment)

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 the product of the potential gift giver is much worse (6 point disadvantage in expected value) the influence of the gift vanishes.[20]

These visual impressions are confirmed by the regression results, reported in Table 2. 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 (BT) and the No Externality (NET) Treatments in column 1, and we add the observations from the Gift

[20] Fligner-Policello robust rank-order tests comparing the average decisions of DMs in the Gift and the No Externality Treatments reveal that we cannot reject the hypothesis that the two are drawn from the same distribution. The *p*-value is .166 (.137) when the gift was given (not given). Subsampling by expected-value differences, the differences in average decisions is significant when the gift was given and *X* has a six point disadvantage (p = .010).

Treatment (GT) in columns 2 and 3. We observe a significant but smaller increase in the frequency of choosing *X* when the gift is sent. Both the increase itself (amounting to 21.0-21.8 percent) and the difference in increase relative to the GT (-12.5 percentage points) are highly significant (p <.01). The decrease after not receiving the gift (-4 percent) is insignificant and the difference in decrease relative to the GT (+8.7 percentage points) is marginally significant (p < .10).[21]

When we account 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 (column 4). The same is true in column 5 when we add the observations from the Gift Treatment, and control for the corresponding interactions "*GT: Gift given*" and "*GT: Gift not given.*" Relative to the effect of gifts in the GT, column 6 shows that there is significantly less punishment for not sending the gift if *X* is the better product and significantly less reward for sending the gift if *X* is much worse.

We draw two main conclusions 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 without externalities. It also explains why many firms offer small gifts to their 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 is much worse than the competing product. For example, when the gift is given, almost 50 percent of all decision makers in the GT choose *X* even though it has a six-point disadvantage after receiving the gift, but less than 5 percent do so in the NET. Thus, the differences in behavior are particularly strong when the external effect is large. This suggests that gift giving can have a large negative impact on social welfare.

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 other decision makers had chosen the prod-

---

[21] Individual-level clustering and bootstrapping methods reach very similar results. With session-level wild cluster bootstrapping, shown in Appendix-Table A4, both the increase after a gift and the difference in increase relative to the GT remain marginally significant even after including controls (p = .087 and p = .088). Note that we do not replicate column 1 of Table 2 in Table A4, because this specification only includes subjects from the BT and the NET, resulting in too few clusters (three) for the wild cluster bootstrap-t procedure to remain accurate.

uct the client preferred. Finally, we asked the clients and producers the same questions (about the decision makers). All subjects were paid for the precision of their estimates using a quadratic scoring rule.

The result is remarkable: In the Gift Treatment decision makers on average estimate 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 66.5 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 over- or underestimate the quality of the decisions.[22] 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.[23]

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

The first set of questions in the Gift Treatment refers to the motivation of the gift giver. When asked why the producers passed on the gift, almost all decision makers strongly agree with the statement that "*the producer wants to influence my behavior*" (1.42 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*

[22] 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 ($p = 0.829$). Wilcoxon-Mann-Whitney tests comparing the decisions of DMs with the predictions of the clients or of the producers, respectively, does not reject the hypothesis that they are drawn from the same distribution ($p = 0.764$ and $p = 0.791$, respectively).

[23] 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 $p = 0.096$. 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 $p = 0.024$.

decision" (4.67). The answers of the clients to these questions were very similar.

The perceptions of DMs and clients closely match the self-reported motivations of producers. Producers openly admit that they offered the gifts "*to influence the decision of the decision maker in my favor*" (1.83), and not "*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 pass on the gift because they want to influence the DM's behavior and because they are afraid that otherwise DMs will not buy their product; DMs and clients perceive this motivation correctly. Nevertheless, DMs respond to the gift.

We also asked DMs directly whether their "*decisions have 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.[24]

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 percent) who admit to having been influenced, 80 percent report positive emotions towards the gift giver or a sense of obligation to buy his product.[25] At the same time, *all* of these 80 percent strongly agree with the statement that the gift was given because "the producer wants to influence my behavior."

6 Gifts that Create Bonds and an Obligation to Reciprocate

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. This contradicts the standard theories of outcome-based social preferences such as altruism, maximin preferences, and inequality aversion, as well as type-based models of reciprocity that all predict that DM favors the client, not the gift giver. Intention-based reciprocity is consistent with the observed behavior, but it is also consistent with the opposite. Furthermore the experiments show that decision makers

[24] 20.8 percent of all DMs strongly agree with the statement "*My decisions have not been influenced*," while 62.5 percent strongly disagree. 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.

[25] When asked whether they "*liked a producer who passed on the gift better than the other producer*," 80 percent of DMs tend to agree. When asked whether the gift giver "*deserves that his product is bought*," 60 percent tend to agree, and when asked whether they "*felt obliged to buy the product*" of the gift giver, again 60 percent tend to agree.

are fully aware that producers give the gift not because they are kind, but because they want to influence their behavior at the expense of their clients. Thus, gift giving is interpreted by the decision makers as a selfish act of the producer.

Why, then, is gift giving so effective? 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 bond between the gift giver and the receiver of the gift. The anthropological and sociological literature is well aware of the fact that a gift creates an obligation that is independent of the intention of the gift giver. In a highly influential essay, the anthropologist Mauss (1924) argues that in archaic societies humans are under an *obligation to give, to receive, and then to repay*.[26] Prominent 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 *internalized*.[27] The view of a gift as a bond that creates an obligation seems to capture the observed behavior better than theories of reciprocity that are based on the type or intention of the gift giver. However, this view of reciprocity has been widely neglected in the economic literature so far.

In the following we propose a simple extension of outcome-based models of social preferences that formalizes this view. The basic idea is that a gift strengthens the bond between the giver and the recipient of the gift, while not giving a gift (if gift giving was possible and expected) weakens it. Suppose that, initially, DM is equally concerned about the payoffs of all other players. If a player increases the decision maker's payoff, that player gets a higher weight in DM's utility function. If a player reduces the decision maker's payoff, the player gets a lower (possibly negative) weight. At a more general level the model endogenizes the "reference group" that each player cares about by making the weight on the material payoff of a player in the utility

[26] Mauss (1924) was inspired by a seminal field study of the Trobrianders (islanders in the Western Pacific) by the anthropologist Malinowski (1922) who identifies 80 forms of social and economic exchange and concludes that all of them are based on reciprocity.

[27] 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. The various synonyms for saying "thank you" reflect this insight: "much obliged" in English, "je vous suis très obligé" in French and "ich bin Ihnen sehr verbunden" (literally: "I am bound to you") in German. The effectiveness of gifts and compliments, even when the recipient is aware of possible ulterior motives, is a building block in the social psychology literature on "ingratiation," which discusses techniques individuals use to bring themselves into the favor of other people (Jones 1964). The marketing literature also shows that 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).

function of another player dependent on the actions taken by the former that affect the latter.[28]

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

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

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:** A player has social preferences over an endogenously formed reference group that can be represented by (2). If player j chooses a pure strategy s_j that increases (decreases) 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 (decreases) as compared to the weight if j had chosen σ_j, i.e.,
>
> $$m^i(s_j, \sigma_{-j}) \geq (\leq) m^i(\sigma_j, \sigma_{-j}) \quad \Rightarrow \quad \alpha_i^j(s_j, \sigma_{-j}|\sigma) \geq (\leq) \alpha_i^j(\sigma|\sigma).$$

Let us apply this simple model to our gift giving game. Suppose that in the Baseline Treatment where no gift can be made the decision maker puts equal weight α on the client and on producers X and Y, i.e., $\alpha_{DM}^C = \alpha_{DM}^X = \alpha_{DM}^Y = \alpha > 0$. This implies that, in the Baseline Treatment, DM favors the client. Consider now the Gift Treatment and the No Externality Treatment, and suppose DM

[28] 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 could 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.

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, $\alpha_{DM}^X(gg\,|\,\sigma_X) > \alpha$, 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* also decreases, $\alpha_{DM}^X(gng\,|\,\sigma_X) < \alpha$, where *gng* indicates that producer *X* did not send the gift ("gift not given.")

In the following proposition, we assume for simplicity that $\alpha_{DM}^X(gg\,|\,\sigma_X) = k \cdot \alpha$, with $0 < \alpha < 1$, and where $k > 1$ is distributed across subjects according to some cdf $F(k)$. Similarly, we assume that $\alpha_{DM}^X(gng\,|\,\sigma_X) = l \cdot \alpha$, where $0 \le l \le 1$ is distributed across subjects according to some cdf $G(l)$. Let $\Delta = m^C(Y) - m^C(X)$ denote the disadvantage of product *X* relative to product *Y* (in terms of expected payoff).

Proposition 4. Consider a decision maker with social preferences satisfying Assumption 1.

(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) worse than *Y*, DM always chooses *Y* in the GT and the NET. If *X* is strictly better, DM may still choose *Y*, and she is more likely to do so when the payoff goes to a third person (GT) than when it goes to herself (NET).

Proof: See Appendix A.

The intuition for these results is straightforward. If producer *X* passes on the gift his weight in DM's utility function increases from α to $k\alpha$. Thus, if *X* is the better product DM chooses *X* in GT and in NET as this choice allows him both to reciprocate towards the gift giver and to benefit her client. If producer *X* passes on the gift but *X* is the worse product, DM chooses *X* in GT if *k* is large enough to offset the utility loss of the client. In NET DM buys the product for herself, so she has a stronger financial incentive to choose the better product. However, if Δ is

small, i.e., product X is not much worse than product Y, and if 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.

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 by which the government requires annual reporting of all physician payments over a cumulative value of USD 100 and which, after September 30, 2014, will be made available to the public.[29] The idea of size limits is to reduce the 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 $420 in a calendar year.[30] The Internal Revenue Service does not allow businesses to tax-deduct a business gift costing more than $25.[31] In Minnesota physicians are not allowed to accept gifts from the pharmaceutical industry that exceed a value of $50 per year.[32] These rules are based on the assumption that large gifts have a large impact on behavior, while the effect of small gifts is small and can be ignored.

We test the effects of such policies in our experimental 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 produc-

---

[29] See p. 9503 of the final rule under www.cms.gov/Regulations-and-Guidance/Legislation/National-Physician-Payment-Transparency-Program/Downloads/Affordable-Care-Act-Section-6002-Final-Rule.pdf , last accessed July 5, 2014.

[30] See http://www.fppc.ca.gov/factsheets/1-09/local_elected.pdf, last accessed July 5, 2014.

[31] See http://www.irs.gov/publications/p463/ch03.html , last accessed July 5, 2014.

[32] See New York Times, 10/12/2007, http://www.nytimes.com/2007/10/12/us/12gift.html, last accessed July 5, 2014.

er 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 given (Panel A) and was not given (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.

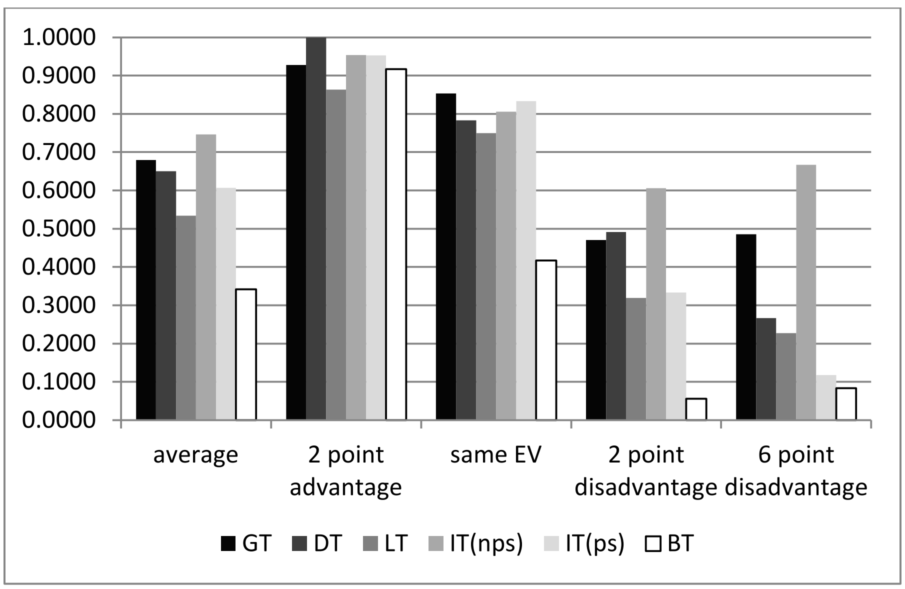

FIGURE 4.A: Frequencies of choosing *X* in the Policy Treatments *when the gift is given* (with comparison to the Baseline and Gift Treatments)

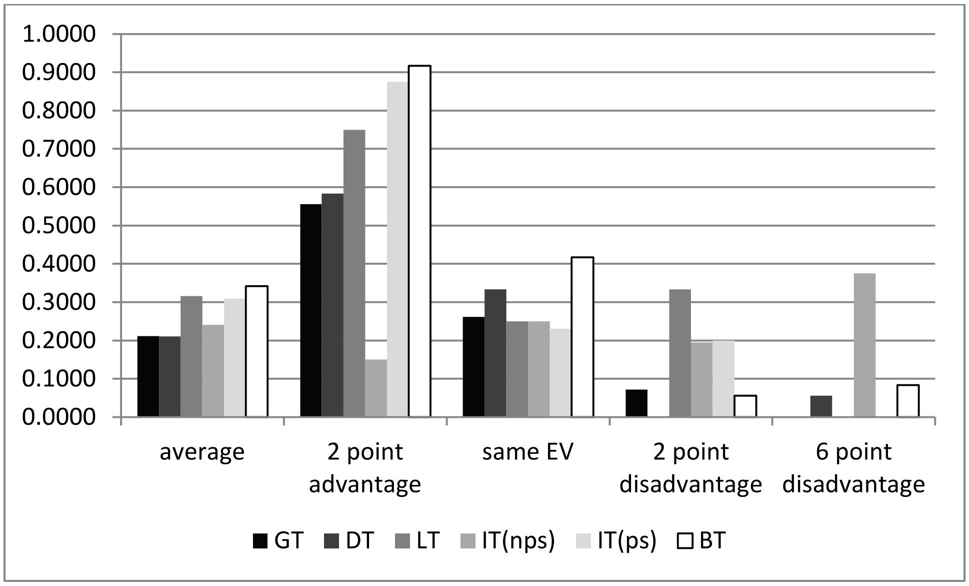

FIGURE 4.B: Frequencies of choosing *X* in the Policy Treatments *when the gift is <u>not</u> given* (with comparison to the Baseline and Gift Treatments)

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 GT (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 but the gift has been sent, decision makers choose the gift giver only in 27 percent of all cases, compared to 49 percent in the GT.[33]

In Table 3, we evaluate the statistical significance of these effects. Columns 1 and 2 replicate the empirical model of Table 2 for the Disclosure Treatment, pooling the observations from the Baseline, Gift, and Disclosure Treatments. The top part of column 1 compares the overall effect of the DT to the Baseline Treatment, mirroring column 2 of Table 2. The bottom part of column 1 compares the effects of DT and BT by subsample, mirroring column 5 of Table 2. 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 2 for the NET). The economically large difference in response to receiving a gift when *X* has a 6 point lower expected value is significant at the five-percent level. This suggests that disclosure may be effective when decision makers do not have

[33] Fligner-Policello robust rank-order tests comparing the average decisions of each DM in the DT when the gift was given (not given) to those in the GT when the gift was given (not given) show that there are no statistically significant differences between these distributions, overall and controlling for expected value differences.

ethical "wiggle room." Still, both these regressions confirm that behavior in the DT is almost identical to the behavior in the original GT without disclosure.[34] In other words, mere "shaming" without actual punishment (such as allowing the client to fire the decision maker) is largely 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 our model. If recipients of a gift reciprocate because the gift creates (or strengthens) a bond between the gift giver and the decision maker, then disclosure should have no effect.

7.2 The Size of the Gift

In the Large Gift Treatment (LT) we triple the size of the gift: if producer X spends one point the decision maker receives six points rather than two points. Six points 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.

Because the effective price of gift giving is reduced it is not surprising that producers send the gift more often. In 92.1 percent of all cases, potential gift givers pass on the gift in the Large Gift Treatment, compared to 71.5 percent in the Gift Treatment. Furthermore, 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 1.83 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 (46 percent).

Surprisingly, however, decision makers respond significantly less to the larger gift, as

[34] Individual-level clustering and bootstrapping methods reach very similar results. With session-level wild cluster bootstrapping, shown in columns 1 and 2 of Appendix-Table A5, the increase after a gift remains significant (p = .007), but the decrease after not receiving a gift becomes insignificant (p = .263).

shown in Figures 4.A and 4.B as well as in columns 3 and 4 of Table 3.[35] Consider first the case where the gift is passed on. In this case decision makers choose the gift giver in 53 percent of all cases, compared to 68 percent in the GT. The difference in coefficients is highly significant ($p < .01$) (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 is significant if *X* has a six-point disadvantage ($p < .01$) and if *X* has a two-point disadvantage ($p < .05$).The difference between LT and GT in the average effect of not passing on the gift is insignificant.[36] However, if *X* has a six-point disadvantage, the effect of not giving a gift is significantly more negative than in the GT ($p < .05$).

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. Similarly, not giving the gift in the LT is farther from the expected strategy σ_j in the LT than in the GT, and therefore the model predicts that decision makers will react more negatively to not receiving the gift in the LT than in the GT.[37]

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 may 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 10 percent of his profits to the decision maker (without knowing whether the producer offers a gift or not). If the client offers profit

[35] Individual-level clustering and bootstrapping methods reach very similar results. For example, with session-level wild cluster bootstrapping, the increase after a gift in column 3 (top part) of Appendix-Table A5 is significant with $p = .022$, and the difference in increase relative to GT (column 4, top part) is marginally significant with $p = .078$.

[36] A Fligner-Policello robust rank-order test comparing the average decisions of each DM when the gift was given in the GT and the LT shows that these differences are highly significant with $p = .027$.

[37] 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 being influenced admit to feeling bad about 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.

sharing he bears half of the cost, i.e., loses 5 percent of his profits.[38] 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 67.2 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). 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 the top portion of Table 3, column 6.[39] 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 3, 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.

Perhaps surprisingly, we find the exact opposite effect 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. For example, when product *X* has a 6 point disadvantage decision makers still choose the gift giver in 66.7 percent of all cases! The regressions in columns 7 and 8 of Table 3 (lower part) confirm that the increase over the Baseline Treatment, estimated to be 59.2 percentage points, is highly significant, while the increase relative to the Gift Treatment, estimated to be 19.1 percentage points, is significant at the five-percent level.[40]

In summary, if the client offers profit sharing DMs react positively to the gift but less so

---

[38] The client pays 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.

[39] Fligner-Policello robust rank-order tests comparing the average decisions of each DM in the Profit Sharing Treatment when profit sharing was offered and the gift was given (not given) to the average decisions of each DM in the Gift Treatment when the gift was given (not given) (controlling for the advantage/disadvantage of the gift giver) show that there is no statistically significant difference between these distributions. Individual-level clustering and bootstrapping methods reach very similar results. For example, with session-level wild cluster bootstrapping, the increase after a gift (when profit is shared) is significant with $p = .044$, and the difference in increase relative to the GT is not significant, with $p = .412$; see Appendix-Table A5.

[40] Fligner-Policello robust rank-order tests comparing the average decisions of DMs in the Gift Treatment to the average decisions of DM in the Profit Sharing treatment when the client has not offered profit sharing show that these differences are statistically significant if the gift is not given and *X* has a two point advantage ($p = .062$) or if the gift is not given and *X* has a six point disadvantage ($p = .012$). With session-level wild cluster bootstrapping, the increase after a gift in top part of column 7 (i.e., profit not shared) is marginally significant ($p = .056$), and the difference in increase relative to the GT (column 8, top part) is insignificant ($p = .186$).

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. As shown in Appendix A (Proposition 5), our model predicts that the effect of gifts 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. The opposite happens when profit sharing is not offered – the bond is weakened and the weight decreases.[41]

This interpretation is also consistent with the questionnaire evidence. The large majority of subjects confirm that 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 become aware that clients could offer profit sharing, they punish clients for not offering 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

---

[41] Note that DMs punish clients who do not share profits rather harshly when the gift is not given (see 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.

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 weaker, in particular when product X is much worse than product Y. Standard models of social preferences have difficulties to predict the observed behavior which is better described by anthropological and sociological theories arguing that gifts create a bond between the gift giver and the recipient and an obligation to reciprocate. We show that a simple extension of standard models of social preferences captures this effect.

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. 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. A more systematic analysis of possible remedies in the lab and in the field is a fascinating and 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.

Camerer, Colin. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics,* Vol. 90, 414-427

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.

Currie, Janet; Lin, Wanchuan and Juanjuna Meng. 2012. "Social Networks and Externalities from Gift Exchange: Evidence from a Field Experiment." *Mimeo*, Princeton University.

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.

Davidson, Russel and Emmanuel Flachaire (2001), "The Wild Bootstrap, Tamed at Last," unpublished manuscript.

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

Efron, B. 1979. "Bootstrapping Methods: Another Look at the Jackknife." *Annals of Statistics*, Vol. 7, 1-26.

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; Goette, Lorenz and Christian Zehnder. 2009. "A Behavioral Account of the Labor Market: The Role of Fairness Concerns." *The Annual Review of Economics*, Vol. 1, 355–384.

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.

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.

Liu, Regina Y. 1988. "Bootstrap Procedures under Some Non-I.I.D. Models." *Annals of Statistics*, Vol. 16, 1696-708

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

Mammen, E. 1993. "Bootstrap and Wild Bootstrap for High Dimensional Linear Models." *Annals of Statistics,* Vol. 21, 255-285

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.

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.

Webb, Matthew. 2012. "Improved Inference in Difference-in-Differences: Monte Carlo Evidence and a Skills Shortage Reduction Program Application." *Working Paper.*

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.

Table 1. Gift Treatment

	OLS	OLS	Logit	OLS	OLS	Logit
	(1)	(2)	(3)	(4)	(5)	(6)
Gift given	0.338***	0.334***	0.336***			
	(0.040)	(0.040)	(0.040)			
Gift not given	-0.130***	-0.131***	-0.164***			
	(0.047)	(0.047)	(0.058)			
Product *X* has higher EV [+2]				0.500***	0.503***	0.546***
				(0.071)	(0.071)	(0.069)
Product *X* has lower EV [-2]				-0.361***	-0.364***	-0.533***
				(0.065)	(0.066)	(0.086)
Product *X* has lower EV [-6]				-0.333***	-0.333***	-0.429***
				(0.071)	(0.071)	(0.082)
(Product *X* has same EV)*(Gift given)				0.436***	0.433***	0.440***
				(0.068)	(0.069)	(0.053)
(Product *X* has same EV)*(Gift not given)				-0.155*	-0.158*	-0.174*
				(0.090)	(0.089)	(0.094)
(Product *X* has higher EV [+2])*(Gift given)				0.011	0.010	0.040
				(0.051)	(0.052)	(0.175)
(Product *X* has higher EV [+2])*(Gift not give				-0.361***	-0.369***	-0.409***
				(0.104)	(0.106)	(0.070)
(Product *X* has lower EV [-2])*(Gift given)				0.415***	0.413***	0.516***
				(0.057)	(0.057)	(0.068)
(Product *X* has lower EV [-2])*(Gift not given				0.016	0.011	0.056
				(0.048)	(0.047)	(0.194)
(Product *X* has lower EV [-6])*(Gift given)				0.402***	0.393***	0.452***
				(0.072)	(0.072)	(0.073)
(Product *X* has lower EV [-6])*(Gift not given				-0.0833**	-0.0699*	
				(0.040)	(0.041)	
Female		0.050	0.062		0.049	0.077
		(0.036)	(0.043)		(0.032)	(0.053)
Economist		0.0612*	0.0740*		0.0608**	0.1000**
		(0.035)	(0.042)		(0.030)	(0.049)
Period		-0.00497*	-0.00595*		0.002	0.003
		(0.003)	(0.004)		(0.003)	(0.005)
Constant	0.342***	0.351***		0.417***	0.356***	
	(0.031)	(0.050)		(0.059)	(0.068)	
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 (evaluated at the mean) in columns (3) and (6). Robust standard errors are reported. 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 2. No Externality Treatment

	Diff. to BT (1)	*Diff. to BT* (2)	*Diff. to GT* (3)	*Diff. to BT* (4)	*Diff. to BT* (5)	*Diff. to GT* (6)
NET: Gift given	0.218***	0.210***	-0.125***			
	(0.052)	(0.051)	(0.048)			
NET: Gift not given	-0.036	-0.043	0.087*			
	(0.049)	(0.048)	(0.050)			
(Product *X* has same EV)*(NET: Gift given)				0.357***	0.350***	-0.084
				(0.082)	(0.082)	(0.068)
(Product *X* has same EV)*(NET: Gift not given)				-0.045	-0.053	0.104
				(0.096)	(0.095)	(0.100)
(Product *X* has higher EV [+2])*(NET: Gift given)				0.099**	0.091**	0.080**
				(0.043)	(0.043)	(0.035)
(Product *X* has higher EV [+2])*(NET: Gift not given)				-0.080	-0.086	0.280**
				(0.079)	(0.079)	(0.118)
(Product *X* has lower EV [-2])*(NET: Gift given)				0.258***	0.248***	-0.166**
				(0.074)	(0.073)	(0.084)
(Product *X* has lower EV [-2])*(NET: Gift not given)				0.039	0.035	0.023
				(0.048)	(0.049)	(0.056)
(Product *X* has lower EV [-6])*(NET: Gift given)				-0.032	-0.038	-0.433***
				(0.056)	(0.057)	(0.072)
(Product *X* has lower EV [-6])*(NET: Gift not given)				-0.068*	-0.076*	-0.005
				(0.041)	(0.041)	(0.010)
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. Robust standard errors are reported. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

Table 3. 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*
	(0.049)	(0.045)	(0.049)	(0.044)	(0.061)	(0.058)	(0.042)	(0.038)
Policy Treatment: Gift not given	-0.120*	0.014	-0.077	0.065	-0.033	0.097	-0.097*	0.026
	(0.064)	(0.065)	(0.111)	(0.112)	(0.077)	(0.078)	(0.052)	(0.054)
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
	(0.080)	(0.065)	(0.083)	(0.067)	(0.091)	(0.078)	(0.076)	(0.060)
(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
	(0.152)	(0.155)	(0.163)	(0.166)	(0.134)	(0.138)	(0.103)	(0.109)
(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)	(0.037)	(0.069)	(0.065)	(0.063)	(0.058)	(0.052)	(0.046)
(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***
	(0.156)	(0.179)	(0.206)	(0.224)	(0.128)	(0.155)	(0.090)	(0.125)
(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*
	(0.072)	(0.083)	(0.062)	(0.074)	(0.109)	(0.117)	(0.066)	(0.078)
(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
	(0.030)	(0.039)	(0.288)	(0.289)	(0.109)	(0.113)	(0.073)	(0.079)
(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**
	(0.090)	(0.101)	(0.075)	(0.087)	(0.088)	(0.099)	(0.081)	(0.093)
(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***
	(0.066)	(0.054)	(0.050)	(0.032)	(0.041)	(0.011)	(0.107)	(0.099)
Dummies for EV differences	X	X	X	X	X	X	X	X
Dummies for EV differences interacted with (GT: gg) and (G	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. Robust standard errors are reported. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

Appendix A [For Online Publication]

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 $(m^X + m^Y + m^C)$, i.e., $\frac{\partial U^{DM}(m^{DM}, m^X, m^Y, m^C)}{\partial(m^X + m^Y + m^C)} > 0$.

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

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

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.[1]

(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.,

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

where $0 < \beta \leq \alpha$.[2] 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. 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: 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

[1] 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.

[2] An alternative formulation is that she wants her own payoff to be as close as possible to the average payoff of all players (Bolton and Ockenfels, 2000), i.e., $U^{DM} = U^{DM}(m^{DM}, m^{DM}/\Sigma_j m^j)$ with $j \in \{DM, X, Y, C\}$ and where $\partial U^{DM}/\partial m^j > 0$ iff $m^{DM}/\Sigma_j m^j > 1/4$. This does not affect our results.

outcome-based preference models predict that DM maximizes the payoff of the client.[3]

(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 as in (a).

(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.***

Proof of Proposition 2: In a pooling equilibrium all producers choose the same strategy. Thus, on the equilibrium path the DM does not learn anything about the type of the (potential) gift giver. If DM is kind she maximizes the sum of payoffs since, in a pooling equilibrium, all other players have the same expected weight in her utility function. Hence, a kind DM favors the client. If DM is selfish, she is indifferent and, hence, also favors the client in equilibrium, as assumed above.

For the proof of the second part of the proposition, 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

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

Thus, any equilibrium in which the selfish type keeps the gift with positive probability must have $p^{gg} - p^{gng} \leq 1/16$. Thus, 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. ***Q.E.D.***

Proof of Proposition 3: At stage 1 producer *X* chooses whether to send the gift (*G*) or not to send the gift (*N*). Then DM decides whether to choose *X*'s product (*X*) or the product that yields the

[3] With altruistic and inequality averse preferences this prediction is unique. If DM has maximin preferences he is indifferent which product to choose.

highest expected payoff for the client (*C*). Since the role of the gift giver was randomly allocated, we maintain that the parties assign a 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 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 $c_{iji} = a_i*$.[4]

If X chooses G then DM always gets 22. If X chooses N then she always gets 20. The equitable payoff for DM, $m^{DM}_{e_X}(b_{X,DM})$, 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."

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 $U^{DM}(X, G, XC) = 22 + n_{DM} \cdot (16-12) \cdot (22-21) = 22 + 4n_{DM}$; by choosing C, she obtains $U^{DM}(C, G, XC) = 22 + n_{DM} \cdot (8-12) \cdot (22-21) = 22-4n_{DM}$. 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, producer 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 $U^X(G, XC, G) = 16 + n_X \cdot (22-21) \cdot (16-12) = 16 + 4n_X$. If she does not pass on the gift she gets $U^X(N, XC, G) = 9 + n_X \cdot (20-21) \cdot (16-12) = 9 - 4n_X$. 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 $U^{DM}(C, N, XC) = 20 + n_{DM} \cdot (9-13) \cdot (20-21) = 20 + 4n_{DM}$; by choosing X, she obtains $U^{DM}(X, N, XC) = 20 + n_{DM} \cdot (17-13) \cdot (20-21) = 22-4n_{DM}$. 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 $U^X(N, XC, N) = 9 +$

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

$n_X \cdot (20\text{-}21) \cdot (9\text{-}13) = 9 + 4n_X$. If she passes on the gift she gets $U^X(G, XC, N) = 17 + n_X \cdot (22\text{-}21) \cdot (9\text{-}13) = 17\text{-}4n_X$. 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

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

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

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

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

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

Hence, if product X is strictly worse ($\Delta > 0$), DM still chooses X for large enough k, and she is more likely to do so in GT than in NET since the 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

$$20+l\cdot\alpha\cdot 16+\alpha\cdot 0+\alpha\cdot m^C(X) \geq 20+l\cdot\alpha\cdot 0+\alpha\cdot 16+\alpha\cdot[m^C(X)+\Delta]$$

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

if X is the (weakly) better product ($\Delta \leq 0$). The inequality cannot hold if X is the worse product ($\Delta > 0$) because $0 \leq l \leq 1$.

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

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

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

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 $0 \leq l \leq 1$, the inequalities hold for large

enough l 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 decision makers with social preferences satisfying Assumption 1.

(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 IT. If product *X* is strictly worse, DM may still choose *X*, and she is most likely to do so in IT with<u>out</u> profit sharing, less likely to do so in GT, and least likely to do so in IT with profit sharing.

(ii) Suppose that producer *X* did not pass on the gift. If *X* is (weakly) worse than *Y*, DM always chooses *Y* in the GT and the IT. If *X* is strictly better, DM may still choose *Y*, and she is most likely to do so in IT with profit sharing, less likely to do so in GT, and least likely to do son IT without profit sharing.

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

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

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

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

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 (Δ > 0), both inequalities still hold for large enough k, small enough Δ, small enough α, or small enough l. 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 k:

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

which always holds. 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 with<u>out</u> profit-sharing than in the GT, we compare again thresholds and note that $1+l\cdot\frac{\Delta}{16}\le 1+\frac{\Delta}{16}$, with the inequality holding strictly for $l<1$.

(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

$$20+0.1\cdot m^C(X)+l\cdot\alpha\cdot 16+\alpha\cdot 0+k\cdot\alpha\cdot 0.95\cdot m^C(X)$$
$$>20+0.1\cdot(m^C(X)+\Delta)+l\cdot\alpha\cdot 0+\alpha\cdot 16+k\cdot\alpha\cdot 0.95\cdot(m^C(X)+\Delta)$$
$$\Leftrightarrow \quad l\cdot 16\alpha>0.1\Delta+16\alpha+k\cdot\alpha\cdot 0.95\Delta$$
$$\Leftrightarrow \quad l>1+\frac{0.95\alpha k+0.1}{16\alpha}\Delta.$$

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

$$20+l\cdot\alpha\cdot 16+\alpha\cdot 0+l\cdot\alpha\cdot m^C(X)>20+l\cdot\alpha\cdot 0+\alpha\cdot 16+l\cdot\alpha\cdot(m^C(X)+\Delta)$$
$$\Leftrightarrow \quad l\cdot 16\alpha>16\alpha+l\cdot\alpha\cdot\Delta$$
$$\Leftrightarrow \quad l\cdot(16-\Delta)>16$$
$$\Leftrightarrow \quad l>\frac{16}{16-\Delta}.$$

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 l, large enough Δ, large enough α, or small enough $\bar{k}$, 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 l:

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

which always holds. 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 with<u>out</u> 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 [For Online Publication]

Period	Po-tential 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.

Appendix C: The Wild Cluster Bootstrap-T [For Online Publication]

A common concern in the analysis of experimental data is the possibility that observations may be correlated within session, due to day effects, (time-varying) experimenter effects or other unobserved factors. One way to address this concern is to allow the error term to be clustered by session. However, standard asymptotic tests tend to over-reject if the number of clusters (i.e., sessions) is small. Therefore, we implement the wild cluster bootstrap-t procedure, first proposed by Cameron, Gelbach, and Miller (2008) to correct standard errors in estimations with few clusters. The general bootstrap method, as introduced by B. Efron (1979), works by generating pseudo-samples from an original sample and calculating a statistic of interest within each pseudo-sample, and then using the distribution of the statistic of interest to infer the distribution of the original sample statistic.

There are many ways to generate pseudo-samples and therefore there exists a wide variety of bootstrap methods. In order for a bootstrap method to be applicable, the resampling method chosen should reflect the original data generating process (DGP) as closely as possible. Wu (1986) first suggested using a bootstrap method known as the wild bootstrap in order to deal with cases with heteroskedastic errors. Liu (1988) and Mammen (1993) provide theoretical justification for using the wild bootstrap in cases with heteroskedastic errors. Cameron, Miller, and Gelbach (2008) extend the wild bootstrap procedure to cases with clusters. They show that their wild cluster bootstrap-t procedure works well even with few clusters (as few as 5).

In our experiments, we utilize the wild bootstrap-t procedure and cluster by both session and subject. The basic argument proposed by Cameron, Miller, and Gelbach (2008) for using certain bootstrap methods over others when the number of clusters is as follows. They propose that the key in these cases is to bootstrap an asymptotically pivotal statistic, meaning a statistic whose asymptotic distribution does not rely on unknown parameters. This leads to asymptotic refinement, meaning the distribution of the test statistic converges faster to the true distribution than test statistics based on conventional asymptotic theory.

For the following, let T be the test statistic which is being bootstrapped. T has a small-sample c.d.f. denoted $H(t|F) = \Pr[T \leq t|F]$ where F is the true c.d.f. generating the underlying data in the sample. The bootstrap estimate of H *is* $\hat{H}\left(t\middle|\hat{F}\right) = \Pr[T^* \leq t|\hat{F}]$ where $\hat{F}$denotes the c.d.f. used to obtain the bootstrap resamples. Assume T is distributed with mean 0 and variance 1. In our analysis, we bootstrap the Wald statistic so these two conditions are satisfied.

The usual asymptotic approximation is given by:

$$\Pr[T \leq t|F] = \Phi(t) + O(N^{-\frac{1}{2}})$$

where $\Phi(.)$ is the standard normal density and N is the sample size. The Edgeworth Expansion, shown below, gives a more accurate asymptotic approximation.

$$\Pr[T \leq t|F] = \Phi(t) + N^{-\frac{1}{2}}a(t)\emptyset(t) + O(N^{-1})$$

where $\emptyset(t)$ is the standard normal density and $a(t)$ is an even quadratic polynomial with coefficients that depend on the low-order cumulants (or moments) of the underlying data.

Now, the bootstrap version of T is T^*, which has Edgeworth expansion

$$\Pr[T^* \leq t|\hat{F}] = \Phi(t) + N^{-\frac{1}{2}}\hat{a}(t)\emptyset(t) + O_p(N^{-1})$$

where $\hat{F}$ is the empirical distribution function of the sample. If $\hat{a}(t) = a(t) + O_p(N^{-\frac{1}{2}})$, which is often the case, then

$$\Pr[T \leq t|F] = \Pr[T^* \leq t|\hat{F}] + O_p(N^{-1})$$

Therefore, the bootstrap c.d.f. is within $O_p(N^{-1})$ of the unknown true c.d.f. $\Pr[T \leq t|F]$, which is a better approximation than the usual asymptotic approximation using $\Phi(t)$, which is within $O(N^{-\frac{1}{2}})$ of the true cdf.

Why is it important to use an asymptotically pivotal test statistic? Suppose T has mean 0 and variance σ^2. Therefore, T/s (where s is a consistent estimate of the standard error) has mean 0 and variance 1. We may still apply Edgeworth expansions, but now

$$\Pr[T \leq t|F] = \Phi(t/\sigma) + N^{-\frac{1}{2}}b(t/\sigma)\emptyset(t/\sigma) + O(N^{-1})$$

And similarly the bootstrap estimates

$$\Pr[T^* \leq t|\hat{F}] = \Phi(t/\sigma) + N^{-\frac{1}{2}}\hat{b}(t/\sigma)\emptyset(t/\sigma) + O_p(N^{-1})$$

Now, even if $\hat{b}(.) = b(.) + O_p\left(N^{-\frac{1}{2}}\right)$, these functions are evaluated at t/s where usually $s = \sigma + O_p\left(N^{-\frac{1}{2}}\right)$. It follows that

$$\Pr[T \leq t|F] = \Pr[T^* \leq t|\hat{F}] + O_p\left(N^{-\frac{1}{2}}\right)$$

So there is no asymptotic refinement when we bootstrap a non-pivotal statistic, i.e., the bootstrap distribution of the test statistic converges no faster to the true distribution relative to first-order asymptotic theory.

Therefore, if a bootstrap method is to provide asymptotic refinement, one must bootstrap an asymptotically pivotal statistic. Cameron, Miller, and Gelbach (2008) suggest the Wald statistic $w = \frac{\widehat{\beta_1} - \beta_1^0}{s_{\widehat{\beta_1}}}$ where $s_{\widehat{\beta_1}}$ is the standard error of $\widehat{\beta_1}$ and $H_o: \beta_1 = \beta_1^0$ and $H_a: \beta_1 \neq \beta_1^0$. Theoretically, we may use a variety of bootstrap methods in order to provide bootstrap estimates of the Wald statistic. Cameron, Miller, and Gelbach (2008) find that the wild cluster bootstrap-t procedure performs particularly well in practice.

The implemented wild cluster bootstrap-t procedure is shown below. In the steps below, we are bootstrapping the Wald statistic for the coefficient on giving a gift. We are comparing behavior of the Gift Treatment to the Baseline Treatment. The procedures for different treatments and for the Wald statistic when the gift was not given are completely analogous.

1. In original sample, estimate the model

$$Y_i = \beta_0 + \beta_{gift} * GiftGiven_i + \beta_{Nogift} * GiftNotGiven_i + u_i$$

   Form Wald statistic for $H_o: \beta_{gift} = 0$

$$w = \frac{\widehat{\beta_{gift}}}{s_{\widehat{\beta_{gift}}}}$$

   where $s_{\widehat{\beta_{gift}}}$ is obtained using the cluster-robust variance estimator. In STATA this is obtained using the cluster option.

2. Estimate the restricted model $Y_i = \beta_0 + u_i$. Obtain the restricted model $\widehat{\beta_0}^R$ and the associated residuals $\{\widehat{u_1}^R, \ldots, \widehat{u_G}^R\}$ where $\widehat{u_g}^R$ is the vector of residuals obtained from cluster g. Note: existing literature on bootstrap methods advocate the use of bootstrap resampling methods that impose the null hypothesis. Therefore, in all of the reported results, we use a bootstrap method which imposes the null hypothesis. Also, had we included demographic controls as regressors (which we do in other cases) our restricted model would have included terms for these controls. We restrict only the coefficients of interest (gift given and gift not given). One may perform the restricted OLS estimation by restricting both coefficients on the statistics of interest simultaneously, or each individually. We perform both procedures.
3. Do B iterations of the next step. On the b^{th} iteration:
4. (a) Form a pseudo sample of G clusters $(\widehat{Y_1^*}, X_1), \ldots, (\widehat{Y_G^*}, X_G)$ by the following method. For each cluster g = 1,…,G, form either $\widehat{u_g^{R*}} = a_g \widehat{u_g}^R$, where $a_g = 1$ with probability 0.5

and $a_g = -1$with probability 0.5. Note that when using weights of this form, the maximum possible unique resamples is equal to 2^G. In many cases we have as few as five clusters, resulting in 32 unique resamples. Webb (2012) argues that inference can be improved by adding points to the weighting distribution. He proposes a six-point weighting distribution and provides Monte Carlo evidence to support using such a distribution. We implemented the wild cluster bootstrap-t procedure using a six-point weighting distribution. The results did not change significantly.

(b) Form $\widehat{Y_g^*} = \widehat{\beta_0}^R + \widehat{u_g^{R*}}$. If other regressors such as demographic controls were included then the bootstrap sample would be formed by also adding the additional regressors multiplied by the estimates of their coefficients in the restricted regression.

(c) Calculate the Wald test statistic $w_b^* = \frac{\widehat{\beta_{gift,b}}^R}{s_{\widehat{\beta^*_{gift,b}}}}$, where $\widehat{\beta_{gift,b}}^R$and its standard error $s_{\widehat{\beta^*_{gift,b}}}$are obtained from the unrestricted OLS estimation using the b^{th} pseudo-sample, with $s_{\widehat{\beta^*_{gift,b}}}$calculated using cluster-robust standard errors.

5. The p-value is calculated by:

$$p^*(w) = 2 * \min\left(\frac{1}{B}\sum_{b=1}^{B} I(w_b^* > w), \frac{1}{B}\sum_{b=1}^{B} I(w_b^* \leq w)\right)$$

where I(.) is the indicator function.

As an illustrative example, we will present a discussion of the original wild bootstrap procedure, taken from Liu (1988). The following displays how the wild bootstrap can be used in order for accurate estimation in cases with heteroskedasticity. This original wild bootstrap does not take into account clustering, but the extension to clusters does not greatly alter the procedure. Therefore, going through the relatively simple example of Liu (1988) allows for a better understanding of why the procedure developed in Cameron, Miller, and Gelbach (2008) provides consistent estimates of the Wald statistic when allowing for clustered errors. Liu (1988) focuses on a simple regression of the form $Y_i = \beta x_i + \varepsilon_i$, where $x_i's$ are nonzero real numbers, $E(\varepsilon_i) = 0$, $Var(e_i) = \sigma_i^2$, and e_i's are independent. The least squares estimate of β is $\hat{\beta} = \frac{\sum_{i=1}^{n} x_i Y_i}{\sum_{j=1}^{n} x_j^2}$. Therefore, in this simple case with one regressor and no constant, $Var(\hat{\beta}) = \frac{\sum_{i=1}^{n} x_i^2 \sigma_i^2}{(\sum_{j=1}^{n} x_j^2)^2}$. Let $u_i = Y_i - x_i\hat{\beta}$ be the residuals. If one employed a traditional residual bootstrap, then the

bootstrap sample would be $Y_i^* = x_i\hat{\beta} + u_i^*$ where $u_1^*, \ldots, u_n^*$ is a random sample drawn from the empirical d.f based on $(u_1 - \overline{u_n}), \ldots, (u_n - \overline{u_n})$. Let $\widehat{\beta_b}$ denote the least squares estimate based on the bootstrap sample. Then the bootstrap variance is $Var\left(\widehat{\beta_b}\right) = \frac{n^{-1}\sum_{i=1}^{n}(u_1 - \overline{u_n})^2}{(\sum_{j=1}^{n} x_j^2)^2}$ which is equivalent to $\frac{n^{-1}\sum_{i=1}^{n}\sigma_i^2}{(\sum_{j=1}^{n} x_j^2)^2}$ asymptotically. Therefore, this bootstrap procedure does not result in a consistent estimate of the standard error of $\hat{\beta}$ if the error variances are allowed to vary. It is easy to alter the bootstrap procedure in order to achieve a consistent estimate. Instead of drawing a bootstrap sample of residuals from the empirical d.f based on $\{(u_i - \overline{u_n})\}$, draw the bootstrap from the empirical distribution based on $\{(x_i/\overline{(x_n^2)}^{\frac{1}{2}}(u_i - \overline{u_n})\}$, where $\overline{(x_n^2)} = \frac{1}{2}\sum_{i=1}^{n} x_i^2$ and let $\widetilde{\beta_b}$ denote the resulting bootstrap least squares estimator of $\hat{\beta}$. Now,

$$Var\left(\widetilde{\beta_b}\right) = \frac{\sum_{i=1}^{n} x_i^2 u_i^2}{(\sum_{j=1}^{n} x_j^2)^2} - \frac{n^{-1}\sum_{i=1}^{n}(x_i u_i)^2}{(\sum_{j=1}^{n} x_j^2)^2}$$

Which is asymptotically equal to

$$\frac{\sum_{i=1}^{n} x_i^2 \sigma_i^2}{(\sum_{j=1}^{n} x_j^2)^2} + O_p(n^{-\frac{3}{2}})$$

So when we draw our bootstrap residuals from the empirical d.f. described above, we get a consistent estimate of the standard error for $\hat{\beta}$.

In order to implement a bootstrap procedure which results in drawing from the desired empirical d.f, Liu (1988) asserts that the bootstrap sample should be

$Y_1^*, \ldots, Y_n^*$ where $Y_i^* = x_i\hat{\beta} + t_i u_i$ where $E(t_i) = 0$ and $Var(t_i) = 1$.

There are many options for t_i. Davidson and Flachaire (2001) provide theoretical justification and Monte Carlo evidence favoring Rademacher weights. Rademacher weights are such that $t_i = 1$ with $p = \frac{1}{2}$ and $t_i = -1$ with $p = \frac{1}{2}$. In all of our reported bootstrap results, we use the Radamecher weights. However, in unreported results, we implemented the bootstrap procedure using the six-point distribution suggested in Webb (2012). The results are similar. Davidson and Flachaire (2001) provide justification for using the residuals and coefficients from a restricted OLS estimation, and therefore the bootstrap sample in the procedure defined above is

$$\widehat{Y_g^*} = \widehat{\beta_0}^R + \widehat{u_g^{R_*}}.$$

Appendix-Table A2. Summary Statistics by Treatment

	Observations					Demographics							
	Total	Subjects	DMs	Producers	Clients	Gender: Females	Age: 20s	30s	40s-60s	Major: 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%

Appendix-Table A3. Gift Treatment

	(1)	(2)	(4)	(5)
Gift given	0.338**	0.334**		
	(0.031)	(0.031)		
Gift not given	0.130*	-0.131		
	(0.094)	(0.156)		
(Product *X* has same EV)*(Gift given)			0.436**	0.433**
			(0.031)	(0.031)
(Product *X* has same EV)*(Gift not given)			-0.155*	-0.158*
			(0.094)	(0.094)
(Product *X* has higher EV [+2])*(Gift given)			0.011	0.010
			(0.970)	(0.842)
(Product *X* has higher EV [+2])*(Gift not given)			-0.361**	-0.369**
			(0.031)	(0.031)
(Product *X* has lower EV [-2])*(Gift given)			0.415**	0.413**
			(0.031)	(0.031)
(Product *X* has lower EV [-2])*(Gift not given)			0.016	0.011
			(0.969)	(0.969)
(Product *X* has lower EV [-6])*(Gift given)			0.402**	0.393**
			(0.031)	(0.031)
(Product *X* has lower EV [-6])*(Gift not given)			-0.083	-0.070
			Could not be estimated	
Dummies for EV differences			X	X
Controls for gender, major, and period		X		X
Observations	720	720	720	720
(Pseudo) R-squared	0.156	0.164	0.407	0.412

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. P-values estimated using the wild cluster bootstrap t-procedure (clustering by session) are reported. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

Appendix-Table A4. No Externality Treatment

	Diff. to BT	*Diff. to GT*	*Diff. to BT*	*Diff. to GT*
	(1)	(2)	(3)	(4)
NET: Gift given	0.210*	-0.125*		
	(0.087)	(0.087)		
NET: Gift not given	-0.043	0.087		
	(0.552)	(0.260)		
(Product *X* has same EV)*(NET: Gift given)			0.350	-0.084*
			(0.230)	(0.054)
(Product *X* has same EV)*(NET: Gift not given)			-0.053	0.104
			(0.534)	(0.465)
(Product *X* has higher EV [+2])*(NET: Gift given)			0.091	0.080
			(0.431)	(0.197)
(Product *X* has higher EV [+2])*(NET: Gift not given)			-0.086	0.280
			(0.493)	(0.210)
(Product *X* has lower EV [-2])*(NET: Gift given)			0.248***	(0.166)
			(0.007)	(0.118)
(Product *X* has lower EV [-2])*(NET: Gift not given)			0.035	0.023
			(0.715)	(0.813)
(Product *X* has lower EV [-6])*(NET: Gift given)			-0.038	-0.433**
			(0.743)	(0.038)
(Product *X* has lower EV [-6])*(NET: Gift not given)			-0.076	-0.005
			(0.135)	(0.599)
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
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
Sample	NET, GT, BT	NET, GT, BT	NET, GT, BT	NET, GT, BT
Observations	1,040	1,040	1,040	1,040
R-square	0.140	0.140	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. P-values estimated using the wild cluster bootstrap-t procedure (clustering by session) are reported. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

Appendix-Table A5. 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
	(0.007)	(0.662)	(0.024)	(0.072)	(0.044)	(0.425)	(0.056)	(0.186)
Policy Treatment: Gift not given	-0.120	0.014	-0.077	0.065	-0.033	0.097	-0.097	0.026
	(0.263)	(0.742)	(0.780)	(0.227)	(0.969)	(0.412)	(0.200)	(0.704)
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
	(0.007)	(0.167)	(0.024)	(0.103)	(0.055)	(0.602)	(0.056)	(0.814)
(Product *X* has same EV)*(Policy T: Gift not given)	-0.076	0.087	-0.208	-0.728	-0.178	-0.021	-0.163	-0.015
	(0.750)	(0.649)	(0.717)	(0.674)	(0.609)	(0.827)	(0.134)	(0.830)
(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
	(0.007)	(0.214)	(0.150)	(0.212)	(0.477)	(0.769)	(0.508)	(0.652)
(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**
	(0.229)	(0.833)	(0.273)	(0.507)	(0.969)	(0.141)	(0.007)	(0.042)
(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
	(0.070)	(0.728)	(0.024)	(0.119)	(0.117)	(0.210)	(0.033)	(0.184)
(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
	(0.263)	(0.382)	(0.273)	(0.434)	(0.345)	(0.248)	(0.915)	(0.223)
(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*
	(0.040)	(0.087)	(0.258)	(0.087)	(0.886)	(0.129)	(0.049)	(0.061)
(Product *X* has lower EV [-6])*(Policyt T: Gift not given)	-0.010	0.061	-0.144	-0.065	-0.07	0.002	0.297	0.364***
	(0.999)	(0.494)	(0.733)	(0.384)	(0.289)	(0.842)	(0.333)	(0.002)
Dummies for EV differences	X	X	X	X	X	X	X	X
Dummies for EV differences interacted with (GT: gg) and (G	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. P-values estimated using the wild cluster bootstrap-t procedure (clustering by session) are reported. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.