CEDEX

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The University of Nottingham

UNITED KINGDOM · CHINA · MALAYSIA

Discussion Paper No. 2018-04

Eugen Dimant

June 2018

Contagion of Pro- and Anti-Social Behavior Among Peers and the Role of Social Proximity

CeDEx Discussion Paper Series
ISSN 1749 - 3293

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/cedex for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
suzanne.robey@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx

Contagion of Pro- and Anti-Social Behavior Among Peers and the Role of Social Proximity

Eugen Dimant

University of Pennsylvania

Abstract

This paper uses a novel experimental design to study the contagion of pro- and anti-social behavior and the role of social proximity among peers. Across systematic variations thereof, we find that anti-social behavior is generally more contagious than pro-social behavior. Surprisingly, we also find that social proximity amplifies the contagion of anti-social behavior more strongly than the contagion of pro-social behavior, and that anti-social individuals are most susceptible to behavioral contagion of other anti-social peers. These findings paired with the methodological contribution are informative for the design of effective norm-based policy interventions directed at facilitating (reducing) pro- (anti-)social behavior in social and economic environments.

Keywords: Behavioral Contagion, Peer Effects, Anti-Social & Pro-Social Behavior

JEL: C91, D64, D9

1. Introduction

This paper is part of a growing literature on social norms, peer-effects, and conformity that has generated substantial interest in the last two decades. Scholars in economics and

*E-Mail: edimant@sas.upenn.edu. I am greatly indebted to the input I received from Ed Glaeser, Dan Houser, Larry Katz, David Laibson, Ulrich Schmidt, and Wendelin Schnedler on the early version of the experimental design. I am particularly thankful for valuable input from Max Bazerman, Cristina Bicchieri, Gary Bolton, Judd Kessler, Rob Kurzban, Muriel Niederle, and Al Roth during my (visiting) research positions at the Harvard University, Stanford University, the University of Pennsylvania, and the University of Texas at Dallas, respectively. I also want to thank Jim Andreoni, Dan Ariely, Robert Cialdini, Rachel Croson, René Fahr, Uri Gneezy, Rachel Kranton, John List, Rosemarie Nagel, David Rand, Imran Rasul, Arno Riedl, and Tim Salmon for interesting discussions and suggestions. This work has greatly benefited from input at the 2016 London Experimental Workshop, the Norms, Actions and Games Workshop at the Institute for Advanced Study in Toulouse, the 2016 CESifo Venice Summer Institute on the Economics of Philanthropy, the 2017 Norms and Behavioral Change Workshop at Nottingham University, the 2017 Behavioral Science Hub Seminar at the LSE, the 2017 Science for Philanthropy Initiative (SPI) conference at the University of Chicago, as well as numerous other workshops and conferences at UC Berkeley, UC San Diego, and Wharton. As part of my dissertation, this paper was awarded the 2015/2016 prize for the best experimental contribution by the German Society for Experimental Economic Research (GfeW). Financial support by the German Research Foundation (DFG) through the SFB 901 is gratefully acknowledged.

This version: June 1, 2018

psychology have advanced our understanding of peer effects and the underlying mechanisms by capitalizing on comprehensive laboratory and field studies over the past decades, particularly with respect to the study of pro-social behavior, cooperation, and reciprocity (Offerman, 2002; Frey and Meier, 2004; Croson and Shang, 2008; Shang and Croson, 2008; Gächter et al., 2013; Thöni and Gächter, 2015). Peer effects on anti-social behavior have been examined with respect to doping (Gould and Kaplan, 2011), dishonesty (Gino et al., 2009) and theft (Falk and Fischbacher, 2002).[1] Peers have also been found to affect academic gains (Duflo et al., 2011), investment decisions (Bursztyn et al., 2014), littering behavior (Cialdini et al., 1990), productivity at work (Ichino and Maggi, 2000; Falk and Ichino, 2006; Mas and Moretti, 2009), juvenile behavior (Damm and Dustmann, 2014), and charitable giving (Meer, 2011; Smith et al., 2015). Most recently, Andreoni et al. (2017) and Fatas et al. (2018) found compelling evidence for preference conformism in response to peer effects. These insights have enriched our understanding of peer effects across various social and economic environments.

Embedded in a controlled laboratory experiment, our paper extends this existing literature in a number of ways, and addresses three main questions. First, does the contagion of pro-social behavior differ from contagion of anti-social behavior?[2] Second, is the contagion of behavior affected by the social proximity to one's peers, and if so to what extent? Third, does contagion affect anti-social individuals differently than it affects pro-social individuals? The last question is particularly relevant from a policy point of view and allows us to better understand who is most susceptible to behavioral change. We investigate these questions by examining the contagion of pro- and anti-social behavior and the role of social proximity with one's peers in the same setting simultaneously. Capitalizing on a controlled and methodologically new framework allows us to address these dimensions simultaneously and to look for policy-relevant interaction effects. Identifying those who are most susceptible to behavioral change enables efficient design of policy measures, for example, in order to micro-target the right individuals and improve the effectiveness of norm- and peer-based interventions to reduce deviance (Miller and Prentice, 2016; Hallsworth et al., 2017).

This paper also makes a methodological contribution to this area of research by address-

[1] Here, we focus on literature closest to our experiment. (See Dimant (2015) for an extensive discussion).

[2] "Behavioral contagion" refers to the phenomenon in which an agent inadvertently changes the behavior of a recipient through social interaction. More precisely, contagion exists when "a recipient's behavior has changed to become 'more like' that of the actor or initiator in which the actor has not communicated intent to evoke such a change." (Wheeler, 1966, p. 179)

ing well-documented challenges in studying peer effects (e.g. endogeneity and reflection problems, as well as selection effects), in which own and peer behavior are circular and reinforcing and ultimately challenge causality (Manski, 1993, 2000; Angrist, 2014).[3] This paper proposes a novel two-step approach as a solution to these challenges, which can be extended to other peer-effect research both in the lab and in the field.[4]

Our approach to studying behavioral contagion among peers is a variation and extension of the give-or-take donation game first introduced by List (2007) and Bardsley (2008), which allows us to create a level playing field to study the contagion of pro- and anti-social behavior simultaneously. In our setup, participants interact with a real charity and can either benefit (through a donation) or harm (through taking away money that was pre-assigned to a charity) a pro-social institution. Our setting is well-suited for a particularly rich analysis of behavioral change, allowing us to examine these questions along the dimensions of frequency (how often behavior spreads) and magnitude (to what extent behavior spreads). Our mechanism for explaining observed behavioral contagion in the pro- and anti-social domain and the role of social proximity is in line with the conformity model by Akerlof (1997).

The design follows a straightforward procedure: action, peer observation, and revision. We consider behavioral contagion to be a change of behavior in the direction of observed behavior. In this sequential variant of a two-player dictator game, the participant (dictator) is matched with a charity (recipient) and is given the opportunity to give or take money away from a charity. After this initial decision, we randomly divide participants into active and passive participants. Passive players are not able to make any changes to their initial decision, while active participants observe the initial behavior of the latter and decide whether to revise their initial decision. This design choice is a core element of our methodological contribution. It addresses the aforementioned theoretical concerns in studying peer effects (Manski, 1993, 2000; Angrist, 2014) because it creates two categories of participants, those who can revise their decision under the influence of peer effects (active participants) and those who cannot revise their decision (passive participants). In addition

[3]Following the literature, we refer to endogeneity and reflection problems as an issue that "arises when a researcher observing the distribution of behavior in a population tries to infer whether the average behavior in some group influences the behavior of the individuals that comprise the group." (Manski, 1993, p. 532). Selection effects indicate that individuals are rarely allocated randomly to peer groups in real life. Rather, individuals choose their peers, hence avoiding an exogenous variation of peer exposure.

[4]For related approaches see Bandiera et al. (2005) and Mas and Moretti (2009).

to observing behavior of one peer, participants are also able to observe varying degrees of social proximity. In our experiment, social proximity is determined by how similar participants' preferences are for things such as food, sports, and family. It is based on a similarity score calculated from congruent answers to a personality questionnaire taken from a major American dating website, which the participants were given at the beginning of the experiment. Common preferences or interests are often the first step in creating a common identity and successful matching (Hitsch et al., 2010).[5]

Across systematic variations of observable proximity to peers and opportunities to react to peer exposure, we find that behavioral contagion is asymmetric in that anti-social behavior is generally more contagious than pro-social behavior. Surprisingly, we also find that social proximity amplifies the contagion of anti-social behavior more strongly than the contagion of pro-social behavior, and that anti-social individuals are most susceptible to behavioral contagion, especially through socially close peers, both in frequency (prevalence of contagion) and magnitude (extent of contagion). Beyond adding to our general understanding of peer effects, these findings point to an important interaction between social proximity and the type of behavioral contagion. This interaction is highly consequential from a policy perspective in designing behavioral interventions, since it provides a lever to influence the behaviors of different types of individuals. Furthermore, this interaction substantiates Cialdini's (2009) *principle of social proof*: obtaining cues about social norms and appropriateness of pro- and anti-social behavior works differently for pro- and anti-social individuals. One potential explanation is that anti-social behavior carries an inherent ambiguity of acceptability until a norm around it is validated through peers, while pro-social behavior is unconditionally acceptable. Because a norm is a collective belief about expected behavior, behavior is more informative if learned from someone who is like oneself (Bicchieri, 2006).

The rest of this paper is structured as follows: we discuss the experimental design and procedure and derive testable hypotheses in Section 2, present the results in Section 3, discuss the results in Section 4 and conclude in Section 5.

[5]In economics, the role of social information in affecting individual behavior has previously been studied by Bohnet and Frey (1999); Charness et al. (2007); Chen and Li (2009); Benjamin et al. (2010); Eckel and Petrie (2011), among others. Social contagion has also been studied in network theory by applying epidemiological techniques used in disease transmission research (Hill et al., 2010).

2. The Experiment

2.1. Design

Researchers face substantial challenges when studying peer effects, especially in field experiments where allocation to peer groups is often non-random (Manski, 1993, 2000; Angrist, 2014). To account for challenges arising from reflection and endogeneity, our experiment centers on two key features: first, after participants make a first give-or-take donation decision, exactly two participants in each session are randomly assigned the role of a passive participant for the remainder of the experiment. In our context, *passive* means that the behavior of these participants is held fixed and is observed by the remaining *active* participants. Assignment to either role was announced after the initial decision. The passive participants do not make any further decisions beyond this point. Holding the observer's behavior constant allows us to control for the aforementioned reflection problem and remove strategic considerations (such as the need to factor in the subsequent behavioral change of the observees). The second key feature is that active participants observe the behavior of the passive participants and can then revise their initial decision. The experiment ends at this point. To ensure that a revision decision can be solely attributed to behavioral contagion, active participants' revision decision remains unobserved.[6]

Another important element of our experiment is the use of social proximity. We mimic social proximity through similar interests and preferences, as taken by a major American dating website. For our purposes, we capitalize on 25 real dating website questions and let the participants in our experiment give binary answers as to whether or not the content of the questions applies to them. Only in the social proximity treatments participants observe both the behavior of the peer and the social proximity among them. In the Baseline, only peer behavior is observable. We discuss the elements of our experiment in detail below.

2.1.1. Basic Experimental Procedure

We implement a two-step give-or-take donation game including a possibility to revise one's initial behavior. Across three treatments, we vary the social proximity signal that

[6]In our study we introduce a defused version of a peer effect study, because the final decision regarding whether or not to revise behavior remains unobservable by one's peers and, thus, carries no social signaling value. Because such a setting distinguishes us from what is typically meant by the term conformity, we use the more encompassing term "behavioral contagion". For a discussion of the mechanism and empirical literature on conformity, see Bernheim (1994). See Wheeler (1966) for a discussion and differentiation from other frequently used terms such as conformity and imitation.

participants observe about their peers. This design allows us to study behavioral contagion as the result of active participants (with non-fixed initial behavior) observing passive participants (with fixed initial behavior) with varying social proximity. To ensure maximum control and randomization in the data collection process, all treatments were played simultaneously within the same experimental session, to which participants were randomly assigned at the beginning of the experiment. The dependent variables of interest are the decision to revise and the exact amount of the revision in response to observing a peer. We test whether and to what extent contagion is driven by both peer behavior and proximity. Figure 1 depicts the experimental design and is explained in detail below.

Consider a dictator game with a taking option in which the participant (dictator) is matched with a charity (recipient) and both parties start with an equal endowment. It is important to note that the endowment given to the charity represents the experimenter's money prepared for donation to the respective charity right after the experiment, which was communicated with the participants. The dictator's action space entails the following: (i) taking away money from the charity, (ii) retaining the distributional status-quo, or (iii) giving money to a charity. Following Eckel and Grossman (1996), we use a charity instead of a participant within each session to increase the saliency of the pro- or anti-social decisions, since the behavior can either benefit or harm a credible institution delivering a public good. For the purpose of this study, we refer to the act of giving money to a charity (at personal cost) or taking money from a charity (for personal gain) as an act of pro-sociality and anti-sociality, respectively.[7]

In the initial stage (Stage 0) participants fill out a personality questionnaire that was later used to determine social proximity among participants. Participants are then sequentially presented with three additional stages, which are the same for all treatments. Importantly, to avoid surprise effects, participants were told at the beginning of the experiment that the experiment consists of multiple stages and that specific information about the nature of the stage will only given once the stages were reached: Stage 1 - action (initial give-or-take decision), Stage 2 - observation (of exactly one passive peer with vary-

[7]The results of our post-experimental questionnaire reveal that participants indeed perceived these decisions to be pro- or anti-social and not simply more or less pro-social. This is in line with the notion of Falk (2017), indicating that our choice paradigm captures not only deviant behavior per se, but also aspects of moral decision-making since the participants can either benefit or harm a pro-social public goods provider (a charity) in an intentional manner, which goes one step beyond studying the contribution (or lack thereof) to a public good.

ing proximity), Stage 3 - reaction (revision of initial give-or-take decision). Once Stage 3 was introduced, the participants were made aware that this concludes the experiment and no further payoff-relevant decisions will be made. The participants then filled out a post-experimental questionnaire and were paid accordingly.

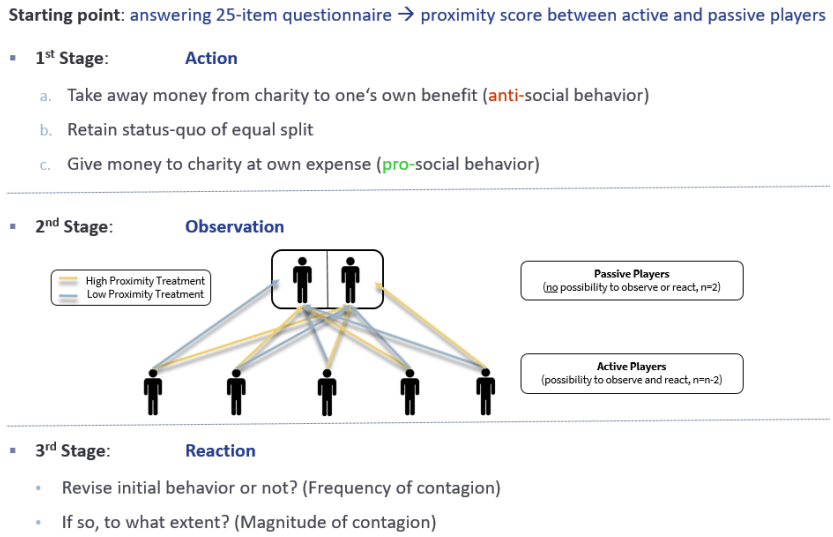

Figure 1: Experimental Design and Procedure. In all treatments, the observed information is the passive player's initial dictator decision. Note that in the Baseline (unknown proximity condition), active players are unaware of the proximity score and only observe behavior.

Stage 0 - Prior to First Decision

Participants answered a personality questionnaire containing twenty-five questions taken from a major US dating website (see the Appendix for details). The answers to these questions were then used to calculate a proximity score among participants in later stages.

Stage 1 - Action

Each participant starts with the same amount of money (€15). Next, participants make one of three mutually exclusive decisions towards a charity[8]: (i) donate their own money

[8]In order to reduce biases and misbehavior stemming from forcing a particular charity on the participants, participants were able to suggest their own charity at the end of the experiment that would receive the respective payment of the one randomly chosen participant. In one case, a participant who was chosen at random could not think of a particular charity. The experimenters then showed the participant a list of the most reputable charities according to CharityWatch.com to choose a charity from their list.

to the charity's account, (ii) maintain the initial equal distribution, or (iii) take money from (and hence harm) the charity and keep the amount for themselves. In order to retain clean observation of initial behavior, information about the specifics of the design were only provided when necessary to reach a deliberate decision. Thus, in Stage 1 participants were neither aware that they might observe peer behavior nor that they might have the opportunity to revise their initial decision in later stages.[9]

Stage 2 - Observation

After each participant makes an initial give-or-take dictator decision, two participants with extreme behavior (either giving away their entire endowment or taking the charity's entire endowment)[10] are assigned the *passive role*. The remaining n-2 participants assume the *active role* and are randomly allocated into one of three treatments. Active player treatments vary by observed behavior and observed social proximity. Such participants are able to observe the behavior of exactly *one* of the two passive peers, while no additional information is given about the behavior of the other, unobserved passive peer. In addition, inferences about the behavior of any other participant, active or passive, are not possible by design because no other information was communicated. In all treatments, the observation of a passive peer contains two pieces of information: the type of behavior (taking or giving) and the exact monetary amount. Treatment variations include the extent to which social proximity information is provided to active participants. They are either given: (i) no information on proximity (Baseline condition) and will observe the behavior either a high or low proximity peer with equal probability without knowing which, (ii) high proximity information, or (iii) low proximity information. More specifically, active participants in the proximity treatments (in addition to observing one passive participant's initial dictator decision of giving to or taking from a charity) are told whether they have either more or less

[9]Importantly, however, in order to make Stage 1 and Stage 3 decisions as comparable as possible, in Stage 1 participants were truthfully told that their decisions will remain anonymous in the sense that neither the experimenter nor other participants will be able to link their give-or-take decision to their identity. More recently, Bolton et al. (2018) show that in the same give-or-take setting the observation through third-parties may in fact amplify the anti-social behavior of the observee.

[10]There were always multiple participants in each session and exactly two participants were randomly chosen. We resort to this approach to keep the observations constant across treatments and sessions and avoiding changes in reference points. In combination with our experimental procedure of having all participants of the same session randomly play any of the three treatments, we are able to retain control over the observation of anti-/pro-social behavior. In addition, this design choice increases statistical power.

similarity to the observed passive peer compared to the second, unobserved passive peer. At no point does the active player receive information about the exact proximity score with the observed passive player. Instead, the active player receives information about the relative proximity score, indicating that the proximity score with the observed passive player is higher or lower than with the second (but unobserved) passive player.

The implementation of the low- and high-proximity information follows a very straightforward calculation. For each participant of the active group, an individual proximity score to both participants of the passive group is calculated based on congruent answers in the list of statements. Across all treatments, the unique advantage of this design is that the same passive participant can be either high or low proximity to active participants:

- **High Proximity Treatment:** Active participants are told that the calculated proximity score (based on the twenty-five item personal questionnaire) between them (the observer) and the observed passive participant is *higher* compared to the second unobserved passive participant. That is, in addition to observing the actual behavior of the passive participant, the active participant is informed that the proximity score is higher with the observed compared to the unobserved passive participant.

- **Low Proximity Treatment:** Active participants are told that the calculated proximity score (based on the twenty-five item personal questionnaire) between them (the observer) and the observed passive participant is *lower* compared to the second unobserved passive participant. That is, in addition to observing the actual behavior of the passive participant, the active participant is informed that the proximity score is lower with the observed compared to the unobserved passive participant.

- **Unknown Proximity Treatment:** The active participant randomly observes the actual behavior of one of the two passive participants, but does not receive any information about the actual proximity scores with either observee.

Stage 3 - Reaction

After observation, the active player is given the option to revise the initial dictator decision. If the active player choses to revise, the revision decision replaces the initial decision. The experiment ends after this stage. It is common knowledge that no other participant would neither be able to observe the revision decision nor the amount revised.

2.1.2. Payoff Structure

Motivated by the aforementioned literature that identifies reflection problems, selection, and strategic interaction as challenges for peer effect studies, in our design the participant's decisions only affected one's own and the chosen charity's payments, rather than those of other participants. Removing strategic interaction among participants ensured that changes in behavior were the direct result of behavioral contagion and not of other-regarding concerns or preemptive behavior (as is the case in Public Goods games, for example). We took additional precautionary steps by designing a payoff mechanism that would cope with potential strategic substitution effects by introducing a 'pay one' random payment mechanism: given the action space in our setup, participants could expect to receive anything between €0 and €20, depending on their choices. However, at the beginning of the experiment it was clearly communicated to the participants that only one participant will be randomly chosen at the end and only his/her decisions will be payoff relevant and determine their own and a charity's payoff. If not chosen randomly, participants would receive their initial endowment (equivalent to €10 per hour), regardless of their decisions towards the charity. This incentive scheme is in line with suggestions made by Charness et al. (2016) and retains incentive compatibility as theoretically argued by Healy et al. (2016).

2.2. Theoretical Arguments and Hypotheses

We nest our testable hypotheses within a simple two-period extension of the quadratic utility conformism model as introduced by Akerlof (1997), which is motivated by the classic work on the economics of conformism by Jones (1984).[11]

For our purposes, we extend Akerlof's reduced form conformity model to derive predictions for behavioral contagion through peers, which also takes into account our solution to the aforementioned reflection problem in studying peer effects. We present an intuition

[11] In particular, we focus on the quadratic utility version of the conformity model by Akerlof (1997) because it generates unique equilibria from which we can derive testable hypotheses for our experiment and it is more in line with the general story of our paper. While Akerlof's general model puts more emphasis on one's own decisions, the quadratic utility model has an a priori assumption that puts one's own and peer behavior on an equal footing with respect to how behavior affects individual utility. Assuming that individuals are on a continuum between extremely selfish and extremely altruistic, the quadratic utility approach is conclusive. It should be noted that this assumption is not crucial to our model's predictions because we mainly focus on the relevance of two factors, i.e. social proximity and the pro-/anti-sociality of observed behavior, in affecting behavioral contagion. Thus, the equal-weight assumption is not decisive in predicting the direction in which individual behavior changes as a function of those two factors.

below and discuss the model and proofs in more detail in the Appendix. Our hypotheses focus on (i) testing the symmetry of behavioral contagion, i.e. magnitudes of contagion in anti-social versus pro-social peer behavior, and (ii) the effect of social proximity with the observed peers on behavioral contagion. Insights from existing experimental research will substantiate our hypotheses.

At the heart of the model is the introduction of a social proximity parameter $\alpha_t^{ij} \in (0,1)$, which depicts the social proximity of individual i with a peer j at time $t \in [1,2]$, which represents the two points in time at which the active participant makes a decision. The proximity parameter moderates the trade-off between an individual's inherent preference for pro- or anti-social behavior θ^i and the actual behavior of the active peer x_t^i and the passive peer x_t^j, respectively. Let Ψ_t^{ij} represent individual i's prior (first period) or actual observation (second period) of individual j's pro-/anti-social behavior. We show that the utility maximization problem is of the following form:

$$\max_{x_t^i} U_t^i = I - (1-\alpha_t^{ij})(x_t^i - \theta^i)^2 - \alpha_t^{ij}(x_t^i - \Psi_t^{ij})^2 \tag{1}$$

In the Appendix we show that in our experimental design the solution to this yields testable hypotheses for behavioral contagion and is represented by what we refer to as *Contagion Gap* as follows:

$$|x_2^i - x_1^j| = |\theta^i - \alpha_2^{ij}\theta^i + \alpha_2^{ij}x_1^j - x_1^j| = |\theta^i(1-\alpha_2^{ij}) + x_1^j(\alpha_2^{ij} - 1)| \tag{2}$$

The intuition behind it is straightforward: the extent to which the behavior of active peer i after peer exposure (t_2) converges to the passive peer's j initial behavior (t_1) is largely driven by the strength of the social proximity α_t^{ij} between peers. The further away individual i's inherent preference for pro-/anti-social behavior is from individual j's initial behavior, the stronger α_t^{ij} has to be in order to close the contagion gap. From this, we can generate our testable hypotheses:

H_1: *Exposure to peer behavior will generate substantial behavioral contagion. The extent of behavioral contagion is symmetric in the pro- and anti-social domain.*

H_2: *Social proximity amplifies both pro- and anti-social behavioral contagion.*

We can also turn to existing experimental research to substantiate our claims and find independent support for our hypotheses. Following the existing literature on social proximity, it is reasonable to assume that observing the behavior of people who are socially closer matters more in terms of what is socially accepted or an existing norm (e.g., Charness et al. (2007)). Existing research indicates that social proximity predicts behavior in different contexts related to charitable giving, trust, punishment, and reciprocity (Akerlof, 1997; Charness et al., 2007; Chen and Li, 2009), as well as neighborhood effects (Damm and Dustmann, 2014). Norms or behavioral prescriptions have also been shown to be associated with one's identity, thus rendering behavioral contagion more likely in situations in which identity to peers is salient (Akerlof and Kranton, 2000; Benjamin et al., 2010). Results on advice seeking and giving substantiate the claim that norm signaling of what is socially acceptable is more relevant coming from people who are socially closer (Gino and Moore, 2007). Finally, from an evolutionary point of view, a number of ultimate-proximate reasons (i.e., kin selection, inclusive fitness, and evolution of fairness) stress the role of proximity and observability of behavior in affecting the likelihood of altruism (Henrich et al., 2010).

3. Results

We conducted the experiment at the BaER-Lab at the University of Paderborn, Germany. Participants were recruited using ORSEE (Greiner, 2015). We used zTree (Fischbacher, 2007) to run our experiment. In total, 227 participants through 9 sessions (ranging from 24 - 26 participants per session) were randomly assigned to one of the three treatments (unknown proximity, high proximity, and low proximity) and we collected data for 83, 68, and 76 participants, respectively. Each session lasted about 45 minutes and the hourly average earnings were roughly $13 (see Appendix Section B Table A1 for details).

We turn to testing our hypotheses by addressing the following three questions: Did behavior spread? And if so, to what extent? And how did proximity affect this contagion, if at all? Our design allows us to study behavioral contagion in two ways: in frequency and in magnitude. The frequency of contagion is a binary measure and equals 1 if, after observing a peer's behavior, the revised behavior converges to the observed behavior. In contrast, for the magnitude of contagion, a value greater (smaller) than zero indicates that the participant is becoming more pro-social (anti-social) relative to one's initial behavior.

In what follows, we first present and analyze the results in Section 3 and reconcile these in a next step with both theoretical and experimental literature in Section 4.

3.1. Behavioral Contagion: Role of Peer Behavior

Our first approach to answering our hypotheses is to look at the frequencies and magnitudes of behavioral contagion conditional on the type of observed behavior. Depending on the active participant's initial behavior, our analysis examines behavioral contagion conditional on observing pro- or anti-social behavior, respectively. Note that our design allows us to do both increase statistical power and retain maximum control over the effect of behavioral contagion because the observed behavior of a passive peer was the same in all sessions and treatments, rather than vary endogenously across treatments.[12]

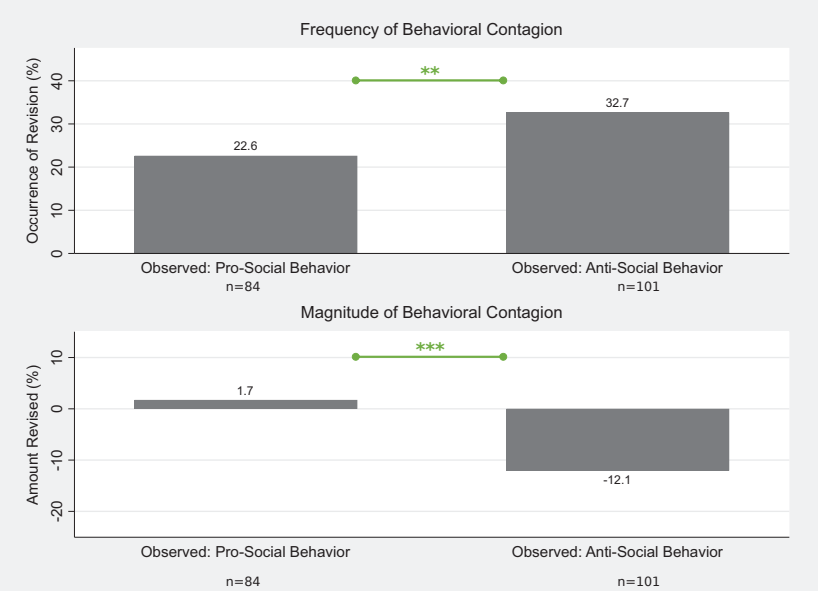

Figure 2: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional. Frequency: prevalence of behavioral contagion (revision of initial behavior into the direction of observed behavior). Magnitude: extent of the revised behavior relative to one's initial behavior. Observed pro-social/anti-social: behavior of the observed passive peer.

First, we look at differences in frequencies of behavioral contagion (upper part of Figure

[12]One alternative way of looking at the data is taking into account the relative difference of own and observed behavior by analyzing participants who observed better behavior (i.e. more pro-social or less anti-social) or worse behavior (i.e. less pro-social or more anti-social) separately. For robustness purposes, we perform such an analysis for a number of cases and present them in Section A of the Appendix. The results are qualitatively congruent with the analysis presented in this section. We introduce controls for the behavioral distance between one's initial and the observed behavior, which is in line with our theoretical framework and the Akerlof (1997) model in our regression analysis.

2). We employ a standard population proportion test and find that observing anti-social behavior triggers behavioral contagion significantly more often than observing pro-social behavior (32.7% vs. 22.6%, p=0.02). Next, we look at the magnitudes of behavioral contagion (lower part of Figure 2). We find that after observing anti-social behavior, participants become 12.1% more anti-social than they were before observing peer behavior. After observing pro-social behavior, participants become more pro-social, but to a lesser extent, by 1.8%. Using Mann-Whitney U (MWU) statistics, we find this difference to be significant at the 1% level (-12.1% vs. 1.7%,p<0.01).[13]

Overall, these results suggest that while exposure to peer behavior indeed leads to substantial contagion, and is thus in line with hypothesis H_1, the extent to which contagion occurs in both behavioral domains is asymmetric. In fact, with a ratio of almost 7:1 (= 12.1/1.7) the downward adjustment of behavior is much more pronounced than the upward adjustment, indicating that anti-social behavior is more contagious than pro-social behavior. This is true in both frequency and particularly in magnitude.

3.2. Behavioral Contagion: Role of Social Proximity

In terms of frequency, we observe contagion significantly more often in the high proximity condition, followed by the low proximity and unknown proximity conditions. Our results, as shown in Figure 3, robustly indicate that higher social proximity indeed triggers stronger behavioral contagion, particularly contagion of anti-social behavior: with increasing social proximity, the frequency and magnitude of revised behavior increases as well. The equality of proportion statistics indicate that the differences are all highly significant. In particular, the occurrence of behavioral contagion in the high proximity condition is more than twice as likely than it is in the low proximity condition (64.9% vs. 29.0%, p<0.01) and more than five times as likely than it is in the unknown proximity condition (64.9% vs. 12.1%, p<0.01). We attribute the latter finding to a saliency effect of observation (Bolton et al., 2018; Bradley et al., 2018) in combination with the fact that participants were matched on similarities rather than dissimilarities, and hence low proximity seems to be perceived as socially closer than unknown proximity. This is also in line with the

[13] An alternative approach to examine differences in magnitudes of behavioral contagion is to follow Moffat Moffatt (2015) and employ the bootstrap two-sample t-test method with 9999 replications. This has the advantage that one can retain the rich cardinal information in the data without making any assumptions about the distribution. Unless noted otherwise, the presented results remain robust and can be obtained from the author upon request.

arguments of our theoretical model.[14] As will be examined next, this significant difference is largely driven by the contagion of anti-social behavior in the high proximity condition.

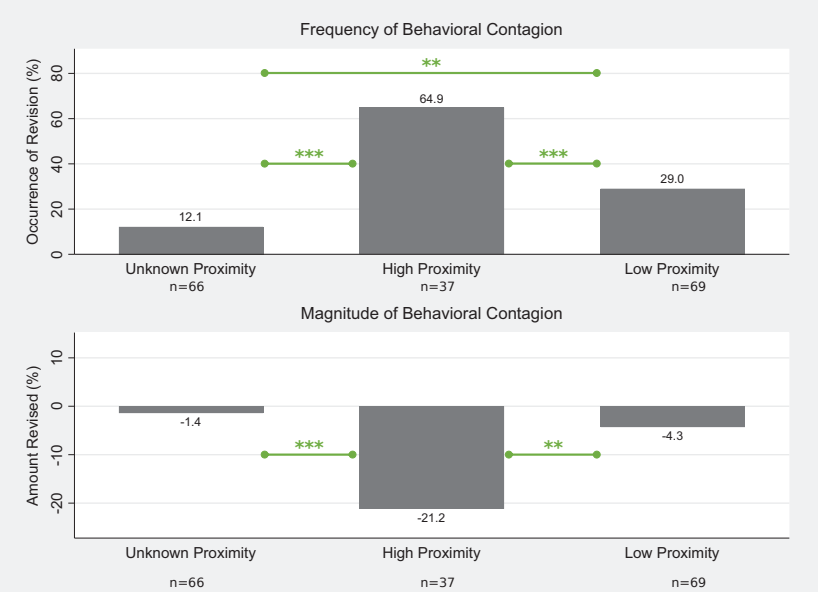

Figure 3: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on social proximity to peer. Definitions of frequency of behavioral contagion and magnitude of behavioral contagion are identical to those used in Figure 2. Unknown / high / low proximity refers to the respective proximity treatment conditions.

Furthermore, we see a similar picture with regards to the magnitude of behavioral contagion. With negative values across all treatment specifications, the results indicate that participants over-proportionally adjust their behavior downwards (becoming more anti-social) than upwards. Again, behavior in the high proximity condition is significantly different from the unknown proximity condition (-21.2% vs. -1.4%, MWU, $p<0.01$) and the low proximity condition (-21.2% vs. -4.3%, MWU, $p=0.02$). We do not observe magnitudes in the low proximity condition to be significantly different from the unknown proximity condition. Our results point to the finding that the salience of low proximity is sufficient to trigger more behavioral contagion in frequency but not magnitude compared to an

[14]Due to the nature of our measure (proximity), these results suggest that providing clues about similarities renders peer behavior more relevant and, thus, leads to more contagion than situations where only the peer's behavior is observed. This is supported by answers to the post-experimental questionnaire. We can only speculate as to whether this effect would be reversed when participants were forced to focus on dissimilarities with their peer, which was not the topic of our project.

environment in which proximity remains unknown. Our subsequent analyses will unpack this finding in more detail. From these results, we conclude that high social proximity triggers significantly higher behavioral contagion in both frequency and magnitude, and is particularly true for the contagion of anti-social behavior. This supports hypothesis H_2.

3.3. Behavioral Contagion: Composite Effect of Peer Behavior and Social Proximity

Next, we combine the results from the previous two subsections and examine the composite effect of observed behavior and social proximity on both the frequency and magnitude of behavioral contagion. The results are presented in Figure 4. For the frequency of contagion, we observe that behavioral contagion is asymmetric only where social proximity is high, with the difference being highly significant according to the equality of proportion statistic (83.3% vs. 40.9%, $p<0.01$). Contagion in the other proximity conditions is less frequent and not significantly different between the types of observed behavior.

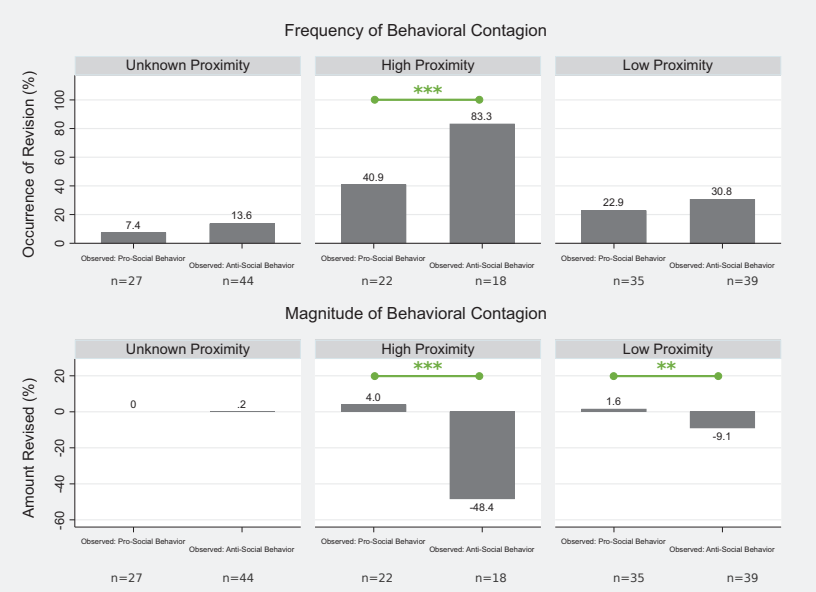

Figure 4: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on social proximity to peer. Definitions of proximity as well as frequency and magnitude of behavioral contagion are identical to those used in previous figures.

The results for the magnitude of behavioral contagion indicate that for both high and low proximity specifications the magnitude of anti-social contagion is larger than the contagion of pro-social behavior, with the test statistics being significantly different at the 1%

level (-48.4 vs. 4.0%, MWU, p<0.01) and 5% level (-9.1% vs. 1.6%, MWU, p=0.011), respectively. Once again, while our results indicate that behavioral contagion exists, we find that the effect is significantly larger and more frequent in the anti-social domain compared to the pro-social domain. In support of hypothesis H_2, we find that social proximity indeed amplifies contagion, particularly for anti-social behavior.

3.4. Whom Does Behavioral Contagion Affect and How?

Lastly, we are interested in understanding whether and to what extent behavioral contagion affects different individuals. Particularly, we will analyze behavioral contagion conditional on the signaled type of the participants, i.e. whether the participants initially engaged in anti-social, pro-social, or equal split behavior. Based on one's initial behavior, our results suggest that behavioral contagion is mostly driven by those who engaged in anti-social behavior, both in frequency and in magnitude. As depicted in Figure 5, in terms of frequency, those participants who decided not to change the initially fair status-quo split between themselves and the charity express behavioral contagion significantly less often than those who initially engaged in anti-social (17.9% vs. 31.3%, p=0.02) or pro-social behavior (17.9% vs. 37.3%, p<0.01). The frequency of contagion in the latter two cases is statistically indistinguishable (31.3% vs. 37.3%, p=0.66).

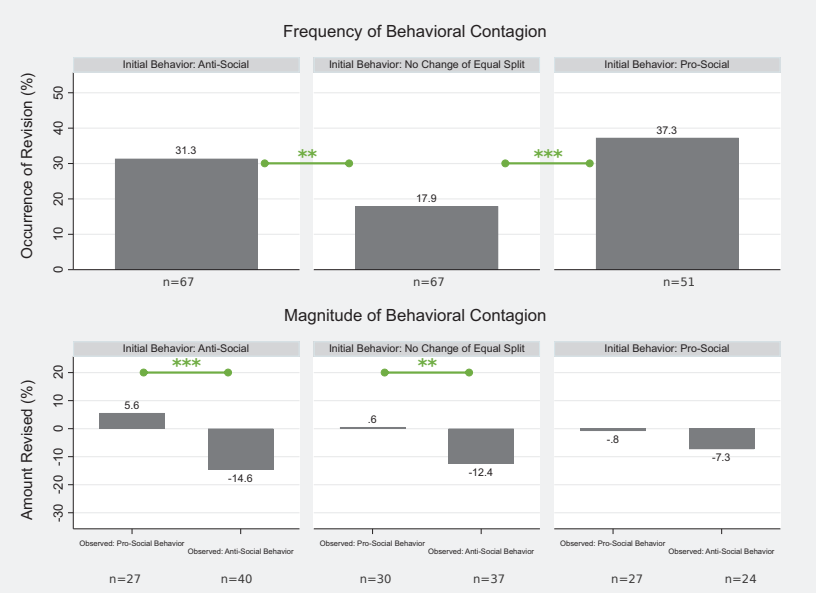

Figure 5: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on initial anti- and pro-social behavior. Definitions of proximity as well as frequency and magnitude of behavioral contagion are identical to those used in previous figures.

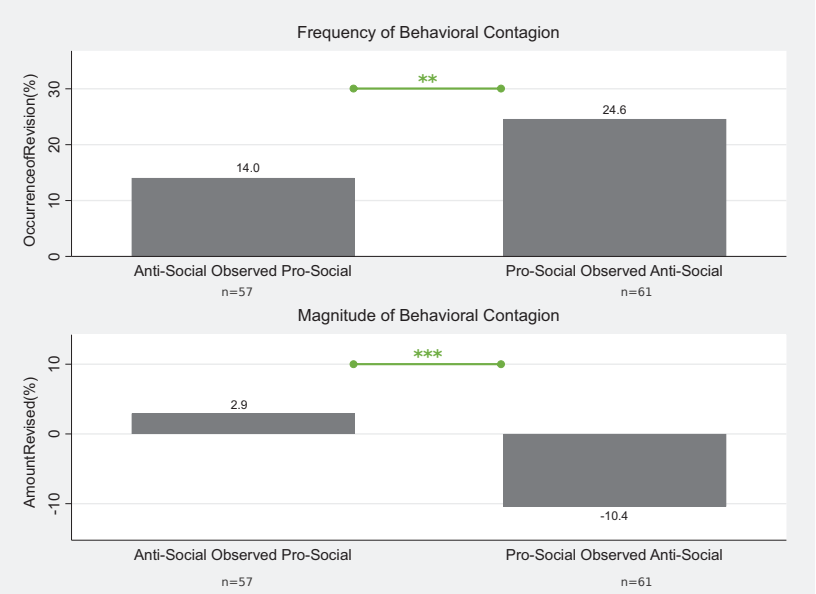

Figure 6: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on observing the opposite behavior than one's own initial behavior. Definitions of proximity as well as frequency and magnitude of behavioral contagion are identical to those used in previous figures.

The pattern is even stronger pronounced when looking at the magnitude of behavioral contagion. When participants initially engaged in anti-social behavior, we observe the largest and most significant asymmetric spread of contagion, indicating that observing more anti-social behavior is almost three times (= |-14.6/5.6|) as contagious as observing pro-social behavior, with the difference being highly statistically significant (-14.6% vs. 5.6%, MWU, $p<0.01$). We observe a similar and highly significant pattern for the initially fair participants (-12.4% vs. 0.6%, MWU, $p=0.03$), but not for those participants who initially engaged in pro-social behavior (-7.3% vs. -0.8%, MWU, $p=0.34$).

In Figure 6, we slice the data even further and specifically look at behavioral contagion of initially anti-social (pro-social) active participants in response to the exposure to passive peers with the opposite, that is pro-social (anti-social) behavior. The results stay qualitatively the same and are significant at the 5% level (24.6% vs. 14.0%, $p=0.02$) and 1% level (-10.4% vs. 2.9%, MWU, $p<0.01$), respectively.

3.5. Regression Analysis

Our previous findings support the notion that anti-social behavior is more contagious than pro-social behavior. Potential explanations include the argument that anti-social behavior carries an inherent ambiguity until the norm is validated through peers, while prosociality is unconditional. Because the norm is a collective belief about expected behavior, behavior is more informative if learned from someone who is like oneself (Bicchieri, 2006). This can explain the observed asymmetry in behavioral contagion and the amplifying effect of social proximity, in particular for anti-social individuals.

In Table 2, we further investigate the robustness of our results using various regression specifications. We analyze the likelihood (Logit) and magnitude (OLS) of contagion with a core set of independent variables that are in accordance with existing theoretical and empirical research (social proximity and the observed anti- or pro-social behavior), an active participant's initial behavioral choice, the initial behavioral gap (amount taken from or given to charity) between oneself and the observed passive participant, prior beliefs about the passive participant's behavior, and gender. We ensure robustness of our results by systematically adding a set of controls, including measures for greed and risk, dummies for the questions used in the personality questionnaire, and relevant interactions that are coherent with our hypotheses.

Our main findings remain robust. Exemplary, we observe that behavioral contagion is both significantly more likely (model 1, 4.4750, $p<0.01$) and also significantly stronger

Table 2: Regression Analysis (Logit and OLS)

	Logit Specifications			**OLS Specifications**		
	DV: Frequency of Contagion (%)			*DV: Magnitude of Contagion (%)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Proximity						
(Base Level: Unknown Proximity)						
High Proximity	30.0168***	17.0782***	7.7620**	-25.2885***	-23.3001***	3.9959
	(21.6982)	(10.4275)	(7.3092)	(8.5234)	(7.8612)	(8.1133)
Low Proximity	4.2097**	3.4353**	2.7482	-4.4984	-4.8457	2.7169
	(2.6540)	(1.7642)	(2.2308)	(4.7669)	(3.9498)	(5.5726)
Observed Anti-Social Behavior	4.4750***	1.4906	3.0349**	-18.4986***	-12.1330*	-1.1459
	(2.4085)	(1.0813)	(1.3537)	(4.9563)	(6.6488)	(5.2042)
Initial Behavior						
(Base Level: No Change of Fair Split)						
Pro-Social	3.7867*	4.9450*	3.1971	8.7674	6.5130	5.2735
	(2.6295)	(4.2757)	(3.0255)	(6.8521)	(8.0419)	(5.4298)
Anti-Social	2.2859	0.3934	1.1170	1.2946	14.9636*	3.9798
	(1.3900)	(0.4201)	(1.1680)	(4.1824)	(7.6048)	(4.0745)
Initial Behavioral Gap	0.9900*	1.0064	0.9884**	0.0143	-0.1056	0.0144
	(0.0058)	(0.0101)	(0.0056)	(0.0758)	(0.0862)	(0.0513)
Beliefs About Peer Behavior						
Pro-Social	0.5854	0.5684	0.6328	-7.2979	-6.5724	-7.3109
	(0.3672)	(0.3320)	(0.3584)	(6.0457)	(5.5534)	(5.6248)
Anti-Social	0.4791	0.4977	0.5450	-4.6446	-2.0336	-1.7580
	(0.2950)	(0.2679)	(0.2927)	(4.7348)	(4.3814)	(4.0170)
Gender	1.1204	0.9249	0.8085	-5.1159	-7.8318	-7.7400*
	(0.7212)	(0.4318)	(0.3681)	(7.4739)	(4.9662)	(4.4730)
Greed		0.9620	0.9903		1.2344	-0.0641
		(0.2088)	(0.2121)		(2.1974)	(1.9323)
Risk		0.8047	0.8595		1.7483	2.2060
		(0.1794)	(0.1895)		(2.2292)	(2.0186)
Interaction 1						
(Initial Behavior x Observed Behavior)						
Pro-Social & Anti-Social		0.5313			4.0753	
		(0.6608)			(12.2119)	
Anti-Social & Anti-Social		23.4464**			-24.4148*	
		(35.8449)			(12.7510)	
Interaction 2						
(Observed Behavior x Social Proximity)						
Anti-Social & High Prox.			6.6283			-56.5932***
			(8.6126)			(14.3517)
Anti-Social & Low Prox.			0.9098			-11.6865
			(0.9670)			(7.6384)
Constant				15.3875	20.1045**	1.4295
				(21.6034)	(9.6465)	(8.2557)
Dummies Questionnaire	Yes	No	No	Yes	No	No
Adjusted R^2				0.089	0.143	0.269
Observations	172	172	172	172	172	172

Notes: Odds ratios (Logit) and regression coefficients (OLS). Standard errors in parentheses. Significances *** 1%; ** 5% level; * 10% level. For Logit, the dependent variable is a dummy that takes on the value 1 if behavioral contagion occurred (= behavior was revised into the direction of the behavior observed). For OLS, the dependent variable is the magnitude in behavioral contagion in percent (= revised amount). 'Proximity' is the degree of social proximity observed by the active participant and takes the value 1 (2) for high (low) proximity. 'Observed anti-social behavior' takes on the value 1 if observed behavior was anti-social. 'Initial behavior' is the active participant's initial decision with value 1 (2) if the participant initially took money away from (gave money to) charity. 'Initial behavioral gap' is the difference between the active participant's initial behavior and the observed behavior (in ECU). 'Beliefs about peer behavior' and 'Gender' are dummies with value 1 if the prior beliefs were correct and for males, respectively. 'Greed' (Hexaco) and 'Risk' (SOEP) are standardized scores where higher values indicate more greediness and higher risk seeking. 'Dummies Questionnaire' are dummies for the items used in the initial personality questionnaire used to generate the proximity score.

(model 4, -18.4986, $p<0.001$) when observing anti-social behavior as compared to pro-social behavior. For the magnitude of contagion, the estimations also indicate that higher peer proximity triggers more behavioral contagion, leading to substantially more anti-social behavior overall (model 4, -25.2885, $p<0.01$). Our results also show that the contagion of anti-social behavior is amplified with higher social proximity. This is true for the impact of high proximity compared to both unknown proximity (model 6, -56.5932, $p<0.001$) and low proximity (model 6, -56.5932 vs. -11.6865, $p<0.01$). In addition, exposure to anti-social peer behavior begets more anti-social behavior, in particular for the initially anti-social individuals. Here, the estimate of Interaction 1 in model 5 indicates that initially anti-social individuals over-proportionally react to anti-social peer behavior compared to their reaction to pro-social peer behavior in comparison to how individuals in the unknown proximity condition react to such peer exposure (model 5, -24.4148, $p=0.057$). A post-estimation test provides suggestive evidence that this result also holds in comparison to how initially pro-social individuals react to such peer exposure (model 5, -24.4148 vs. 4.0753, $p=0.090$). These results also hold with respect to the frequency of behavioral contagion (model 3), in that initially anti-social individuals over-proportionally react to anti-social peers (23.4464, $p=0.039$), also in comparison to the initially pro-social individuals (23.4464 vs. 0.5313, $p=0.070$).

All remaining controls are not robustly significant across our specifications and do not affect the main results. As previously mentioned, we elicited incentivized beliefs (€0.50 upon correct answer) about peer behavior after one's initial decision but prior to peer observation. As the regression results suggest, such beliefs yielded no explanatory power.[15] In sum, the regression results support our non-parametric results presented before, indicating that anti-social behavior is more contagious, which is driven by social-proximity, and that initially anti-social individuals are more susceptible to behavioral contagion, in particular with respect to over-proportionally reacting to anti-social as compared to pro-social peers.

[15] Beliefs about the peer's behavior were elicited after providing information about the treatment condition (high proximity, low proximity, unknown proximity) that the participant is in. A separate non-parametric comparison of these beliefs yield no significant differences, indicating that proximity and anticipation of a peer's initial behavior are unrelated.

4. Discussion

We find conclusive evidence for behavioral contagion and support for our hypotheses, in particular with respect to the role of social proximity among peers. Interestingly, unlike the predictions from our theoretical model and existing literature, we find behavioral contagion to be asymmetric towards anti-social behavior.

In short, our results are both surprising and novel. We can reconcile these findings with existing experimental literature. Pro-social behavior implies bearing costs in order to improve the well-being of others. Anti-social behavior often implies improving one's own well-being at the expense of a third party. This reasoning is coherent with findings on the asymmetry between positive and negative reciprocity (Offerman, 2002), the increase and decrease of pro-social behavior in contributing to a public good (Croson and Shang, 2008; Chen et al., 2010) and the findings on slippery-slope effects (Gino and Bazerman, 2009). This assumption is also in line with the recent findings of Smerdon et al. (2016) and Gächter et al. (2018) on the role of bad behavior and the perseverance of bad norms. In addition, research also indicates that bad social norms can be particularly sticky and thus likely favor the asymmetry towards the contagion of anti-social behavior, as observed in our experiment (Kling et al., 2005; Dimant et al., 2015; Smerdon et al., 2016). Results by Charness et al. (2017) who theoretically and experimentally examine conformism in strategic sequential trust games are also in support of this finding and substantiate our results. Among other things, they find that individuals are more likely to imitate actions that are more self-interested than their own initial choice. Another relevant finding in the literature is that anti-social behavior carries an inherent ambiguity unless the norm is validated, i.e., via observation, while pro-social behavior is unconditional (Bicchieri, 2006), which speaks to the asymmetry in contagion since peer effects affect one's perception of norms. Moreover, our findings on the magnifying effect of social proximity are in line with the findings of Charness et al. (2007), indicating that a salient state of social identity triggers favoritism towards those of stronger social kinship. If we follow the idea that the participants are often uncertain about what the norm is, especially when facing unethical or anti-social dilemmas (Bicchieri and Dimant, 2018), then the observation of peer behavior might receive more weight when it is from a similar other. Because a norm is a collective belief about expected behavior, behavior is more informative if learned from someone who is like oneself (Bicchieri, 2006). This, in turn, can explain our results and the magnified

impact of social proximity on contagion.[16]

In sum, our results highlight the importance of our contribution in this paper: peer effects are not uniform, but rather strongly dependent on both the (anti-)sociality of observed behavior and the degree of social proximity to the observed peer. We treat these insights as novel in that they provide a better understanding of peer effects and behavioral contagion. In real life, the appropriateness of anti-social behavior is thought to be more ambiguous than that of pro-social behavior, which is usually unambiguously appropriate. The nature of anti-social behavior implies the overstepping of social boundaries or the violation of laws. It can thus be expected that individuals often might want to engage in anti-social behavior, but seek social cues and signals to justify engaging in such behavior themselves. After all, a crook among crooks is likely to be more appreciated than a crook among do-gooders. Our interpretation of these findings is that, in comparison to one's peers, individuals look for (mental) excuses to be less selfish, but are more eager to embrace a peer signal that it is socially acceptable to be selfish. These finding suggest that the initially anti-social individuals are most susceptible to anti-social peer behavior in morally loaded environments, which substantiates Cialdini's (2009) *principle of social proof*.

5. Conclusion and Outlook

Deviant behavior that benefits oneself at the expense of others is socially harmful. From a welfare perspective, it also results in distortive, second-best solutions. Conversely, the voluntary redistribution of money to those who have the least, e.g. in the form of donations, is socially desirable. This renders it worthwhile to understand the underlying mechanisms that drive the contagion of pro- and anti-social behavior. Understanding how peer effects are affected by social interactions helps to better understand societal and

[16] We address potential experimenter demand concerns with transparency during the experiment and do not withholding information from the participant with regards to either the matching algorithm or the treatment variations. That is, participants are told that there is a random but equal chance that they would end up in the unknown, low, or high proximity treatment after being told about the details of all potential treatments. This transparency in combination with the fact that all treatments were played within the same session is a strong argument against confounds stemming from demand effects. Arguably, this approach enables participants to interpret the presented social proximity information as the result of chance, rather than as additional information willingly implemented by the experimenter to trigger any particular behavior. Comments made by Ed Glaeser, Dan Houser, Larry Katz, David Laibson, Ulrich Schmidt, and Wendelin Schnedler on an early version of this experiment helped to implement the design choices accounting for EDE as described in this chapter (for a discussion, see Zizzo (2010)).

economic outcomes beyond what standard economic forces can explain. It also allows us to implement better-targeted policy measures for a battery of challenges such as reducing crime rates, improving health conditions, and increasing labor market participation.

Although peer effects have been extensively studied in different contexts both in the field and in the lab, this paper aims to improve our understanding of whether, to which extent, and through which channels individuals are influenced by their peers to engage in more pro- or anti-social behavior. A wealth of literature points at methodological challenges in studying peer effects, including engodeneity, self-selection, and reflection problems (Manski, 1993, 2000; Angrist, 2014). We capitalize on a novel laboratory design that allows us to study behavioral contagion among peers along different dimensions of pro- and anti-sociality as well as social proximity. Our design proposes a solution for the aforementioned methodological challenges. We introduce a variant of a give-or-take donation game (List, 2007; Bardsley, 2008), in which participants are able to give money to or take money from a charity. Novel features of the design include a revision option following the exposure to peer behavior, as well as active and passive roles in which observed behavior is systematically held fixed. Across treatments we vary the social proximity to the observed peer.

By extending existing research from both the methodological and the content perspective, our work contributes to a better understanding of the nature of peer effects and behavioral contagion. In particular, we answer the following three questions: 1) To what extent does the pro-/anti-sociality of a peer's observed behavior influence one's own behavior? 2) What is the role of social proximity with one's peers in affecting behavioral contagion? and 3) Who is affected by behavioral contagion the most? The latter is particularly important for the design of norm-based policy interventions because it is informative or whose behavior is most susceptible to peer influences and in what form.

Our results suggest that anti-social behavior is more contagious than pro-social behavior, in both frequency and magnitude. As expected and motivated by theoretical and empirical insights, social proximity among peers amplifies behavioral contagion. Surprisingly, however, social proximity disproportionately amplifies the contagion of anti-social behavior. We also find that behavioral contagion affects initially anti-social peers in both directions, pro- and anti-social, which is amplified if contagion arises from socially close peers. For the study of peer-effects our results suggest that individuals are frequently at the tipping point of behaving anti-socially. Interestingly, since the nature of our implemented revision option retained full anonymity and thus contained no signaling value, our findings suggest that behavioral contagion occurs even if it merely contains self-signaling value.

More generally, our findings are particularly consequential from a policy perspective in that they allow the identification of the target group that is most susceptible to behavioral change, which is particularly important for norm-based interventions (Miller and Prentice, 2016). The results also help to better understand the "bad apple" model of peer effects and inform policy interventions in education and work environments. There, isolated initial bad behavior can spread uncontrollably, as can be observed from the almost daily scandals featured in the news. It is important to understand the extent to which these differences are driven by the two factors studied in this paper: the exposure to pro-/anti-sociality of peer behavior and the magnitude of social proximity with one's peers. One concrete policy implication is that in addition to providing information about peer behavior, as is done with norm-based policy instruments (Miller and Prentice, 2016; Hallsworth et al., 2017), making use of social proximity information as well as our insights on who is most susceptible to behavioral contagion can help to direct the influence of peer-effects towards the increase pro-social behavior (e.g., charitable giving) and reduction of anti-social behavior (e.g., smoking). We hope our results and methodological contribution inform more effective norm- and peer-based interventions and facilitate further research.

References

Akerlof, G. A. (1997). Social distance and social decisions. *Econometrica: Journal of the Econometric Society*, pages 1005–1027.

Akerlof, G. A. and Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, 115(3):715–753.

Andreoni, J., Nikiforakis, N., and Siegenthaler, S. (2017). Social change and the conformity trap. Technical report, Mimeo.

Angrist, J. D. (2014). The perils of peer effects. *Labour Economics*, 30:98–108.

Bandiera, O., Barankay, I., and Rasul, I. (2005). Social preferences and the response to incentives: Evidence from personnel data. *The Quarterly Journal of Economics*, 120(3):917–962.

Bardsley, N. (2008). Dictator game giving: altruism or artefact? *Experimental Economics*, 11(2):122–133.

Benjamin, D. J., Choi, J. J., and Strickland, A. J. (2010). Social identity and preferences. *American Economic Review*, 100(4):1913–28.

Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy*, 102(5):841–877.

Bicchieri, C. (2006). *The grammar of society*. Cambridge University Press.

Bicchieri, C. and Dimant, E. (2018). It's not a lie if you believe it: ambiguity and cheating behavior. Technical report, Mimeo.

Bohnet, I. and Frey, B. S. (1999). Social distance and other-regarding behavior in dictator games: Comment. *American Economic Review*, 89(1):335–339.

Bolton, G., Dimant, E., and Schmidt, U. (2018). The dark side of reputation. Technical report, Mimeo.

Bradley, A., Lawrence, C., and Ferguson, E. (2018). Does observability affect prosociality? *Proceedings of the Royal Society B: Biological Sciences*.

Bursztyn, L., Ederer, F., Ferman, B., and Yuchtman, N. (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*, 82(4):1273–1301.

Charness, G., Gneezy, U., and Halladay, B. (2016). Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization*, 131:141–150.

Charness, G., Naef, M., and Sontuoso, A. (2017). Opportunistic conformism. Technical report, University of Pennsylvania Working Paper Series.

Charness, G., Rigotti, L., and Rustichini, A. (2007). Individual behavior and group membership. *American Economic Review*, 97(4):1340–1352.

Chen, Y., Harper, F. M., Konstan, J., and Li, S. X. (2010). Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review*, 100(4):1358–98.

Chen, Y. and Li, S. X. (2009). Group identity and social preferences. *American Economic Review*, 99(1):431–57.

Cialdini, R. B. (2009). *Influence: Science and practice*, volume 4. Pearson Education Boston, MA.

Cialdini, R. B., Reno, R. R., and Kallgren, C. A. (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology*, 58(6):1015.

Croson, R. and Shang, J. Y. (2008). The impact of downward social information on contribution decisions. *Experimental Economics*, 11(3):221–233.

Damm, A. P. and Dustmann, C. (2014). Does growing up in a high crime neighborhood affect youth criminal behavior? *American Economic Review*, 104(6):1806–32.

Dimant, E. (2015). On peer effects: Behavioral contagion of (un) ethical behavior and the role of social identity. Technical report, University Library of Munich, Germany.

Dimant, E., Krieger, T., and Redlin, M. (2015). A crook is a crook… but is he still a crook abroad? on the effect of immigration on destination-country corruption. *German Economic Review*, 16(4):464–489.

Duflo, E., Dupas, P., and Kremer, M. (2011). Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in kenya. *American Economic Review*, 101(5):1739–74.

Eckel, C. C. and Grossman, P. J. (1996). Altruism in anonymous dictator games. *Games and Economic Behavior*, 16(2):181–191.

Eckel, C. C. and Petrie, R. (2011). Face value. *American Economic Review*, 101(4):1497–1513.

Falk, A. (2017). Status inequality, moral disengagement and violence. Technical report, IZA Discussion Papers.

Falk, A. and Fischbacher, U. (2002). "Crime" in the lab-detecting social interaction. *European Economic Review*, 46(4-5):859–869.

Falk, A. and Ichino, A. (2006). Clean evidence on peer effects. *Journal of Labor Economics*, 24(1):39–57.

Fatas, E., Heap, S. P. H., and Arjona, D. R. (2018). Preference conformism: An experiment. *European Economic Review*.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Frey, B. S. and Meier, S. (2004). Social comparisons and pro-social behavior: Testing conditional cooperation in a field experiment. *American Economic Review*, 94(5):1717–1722.

Gächter, S., Molleman, L., and Nosenzo, D. (2018). Peer effects in norm compliance. Technical report, Mimeo.

Gächter, S., Nosenzo, D., and Sefton, M. (2013). Peer effects in pro-social behavior: Social norms or social preferences? *Journal of the European Economic Association*, 11(3):548–573.

Gino, F., Ayal, S., and Ariely, D. (2009). Contagion and differentiation in unethical behavior: The effect of one bad apple on the barrel. *Psychological Science*, 20(3):393–398.

Gino, F. and Bazerman, M. H. (2009). When misconduct goes unnoticed: The acceptability of gradual erosion in others' unethical behavior. *Journal of experimental Social psychology*, 45(4):708–719.

Gino, F. and Moore, D. A. (2007). Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making*, 20(1):21–35.

Gould, E. D. and Kaplan, T. R. (2011). Learning unethical practices from a co-worker: the peer effect of jose canseco. *Labour Economics*, 18(3):338–348.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125.

Hallsworth, M., List, J. A., Metcalfe, R. D., and Vlaev, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of Public Economics*, 148:14–31.

Healy, P. J., Azrieli, Y., Chambers, C. P., et al. (2016). Incentives in experiments: a theoretical analysis. Technical report.

Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., et al. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science*, 327(5972):1480–1484.

Hill, A. L., Rand, D. G., Nowak, M. A., and Christakis, N. A. (2010). Infectious disease modeling of social contagion in networks. *PLOS Computational Biology*, 6(11):e1000968.

Hitsch, G. J., Hortaçsu, A., and Ariely, D. (2010). Matching and sorting in online dating. *American Economic Review*, 100(1):130–63.

Ichino, A. and Maggi, G. (2000). Work environment and individual background: Explaining regional shirking differentials in a large italian firm. *The Quarterly Journal of Economics*, 115(3):1057–1090.

Jones, S. R. (1984). *The economics of conformism*. Blackwell.

Kling, J. R., Ludwig, J., and Katz, L. F. (2005). Neighborhood effects on crime for female and male youth: Evidence from a randomized housing voucher experiment. *The Quarterly Journal of Economics*, 120(1):87–130.

List, J. A. (2007). On the interpretation of giving in dictator games. *Journal of Political Economy*, 115(3):482–493.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542.

Manski, C. F. (2000). Economic analysis of social interactions. *Journal of Economic Perspectives*, 14(3):115–136.

Mas, A. and Moretti, E. (2009). Peers at work. *American Economic Review*, 99(1):112–45.

Meer, J. (2011). Brother, can you spare a dime? peer pressure in charitable solicitation. *Journal of Public Economics*, 95(7-8):926–941.

Miller, D. T. and Prentice, D. A. (2016). Changing norms to change behavior. *Annual Review of Psychology*, 67:339–361.

Moffatt, P. G. (2015). *Experimetrics: Econometrics for experimental economics*. Palgrave Macmillan.

Offerman, T. (2002). Hurting hurts more than helping helps. *European Economic Review*, 46(8):1423–1437.

Shang, J. and Croson, R. (2008). Field experiments in charitable contribution: The impact of social influence on the voluntary provision of public goods. *The Economic Journal*, 119(540):1422–1439.

Smerdon, D., Offerman, T., and Gneezy, U. (2016). Everybody's doing it: On the emergence and persistence of bad social norms. Technical report, Tinbergen Institute Discussion Paper.

Smith, S., Windmeijer, F., and Wright, E. (2015). Peer effects in charitable giving: Evidence from the (running) field. *The Economic Journal*, 125(585):1053–1071.

Thöni, C. and Gächter, S. (2015). Peer effects and social preferences in voluntary cooperation: A theoretical and experimental analysis. *Journal of Economic Psychology*, 48:72–88.

Wheeler, L. (1966). Toward a theory of behavioral contagion. *Psychological Review*, 73(2):179.

Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics*, 13(1):75–98.

Appendix

Section A: Theory

Intuitively, one could reasonably assume that the behavior of socially closer people exhibits a more persuasive and informative influence on one's decision making than the behavior of random strangers. The purpose of this section is to outline a simple reduced-form model that allows us to draw predictions about individual behavioral decisions reflected by the experimental design. The intuition behind Akerlof's approach is that individual decision-makers face a trade-off between their own behavioral preferences and the preferences of peers. The magnitude to which one is willing to revise one's own initial behavior and adapt is first and foremost a function of the social proximity of peers. "As a consequence, the impact of my choices on my interactions with other members of my social network may be the primary determinant of my decision, with the ordinary determinants of choice the direct additions and subtractions from utility due to the choice of only secondary importance" (Akerlof, 1997, p. 1006-1007). As a consequence, the individual's utility is subject to a relative evaluation of her own behavior and the behavior of peers. Similar to Akerlof (1997), individual utility is declining with increasing distance between one's own behavior and the peer's behavior.

We resort to a two-period give-or-take donation game in which each individual is paired with a charity. Both the individual and the charity start with an initial endowment $I_c = I_i = I \in R^+$ of equal size. At each period $t \in [1, 2]$, with t_1 (t_2) corresponding to Stage 1 - Action (Stage 3 - Reaction) in our experiment, each individual i faces the choice, x_t^i, of either (a) donating part or all of her own money to the charity, (b) retaining the equal split, or (c) taking away part or all of the charity's money and adding it to her own income. We will refer to (a) and (c) as pro-social or anti-social behavior, respectively. Naturally, the individual's decision is of the form $x_t^i \in [-I, +I]$. The only difference between both periods is the information set that the individual possesses about her peer's behavior. That is, after completion of period 1, the individual observes a random individual's behavior from period 1. In period 2 (after the observation), the individual is given the opportunity to revise her initial decision, if desired.

Let $\alpha_t^{ij} \in (0, 1)$ depict the social proximity of an individual at time t, and θ^i describe the individual's inherent attitude towards pro-/anti-social behavior. That is, θ^i represents the individual i's preference to give or take a particular monetary amount within the boundaries of one's income in a given situation, thus being defined as $\theta^i \in [-I, +I]$. Let

Ψ_t^{ij} represent individual i's prior (first period) or actual observation (second period) of individual j's pro-/anti-social behavior. In a situation in which no outside information is available (period 1), we assume the individual to proxy her beliefs with her own inherent preference for such behavior.

$$\Psi_t^{ij} = \begin{cases} \theta^i, & \text{if } t = 1 \text{ (belief about other's behavior)} \\ x_1^j, & \text{else (observed behavior)} \end{cases} \tag{1}$$

It follows from this that each player maximizes her own utility in the following form:

$$\max_{x_t^i} U_t^i = I - (1 - \alpha_t^{ij})(x_t^i - \theta^i)^2 - \alpha_t^{ij}(x_t^i - \Psi_t^{ij})^2 \tag{2}$$

In the following sections, we examine behavior in the first and second period separately since participants can revise their initial decision.

First Period

Individual i has no information about individual j's decision as all participants carry out their decisions simultaneously and anonymously. Consequently, in order to inform her own initial decision, i resorts to forming beliefs about the behavior of j and, for the sake of simplicity, assumes j's behavior to equal her own inherent attitude towards pro- and anti-social behavior. Proper experimental randomization ensures an equal distribution of beliefs across treatments. The maximization problem is reduced to the following assessment:

$$\frac{\partial U_t^i}{\partial x_t^i} = 2\alpha_t^{ij}(\theta^i - x_t^i) - 2(\alpha_t^{ij} - 1)(\theta^i - x_t^i) \Leftrightarrow x^* = x_t^i = \theta^i \tag{3}$$

We conclude that under anonymity (period 1), the individual's decision to engage in pro- or anti-social behavior is simply the result of her inherent characteristics.

Second Period

Active players receive a signal about a peer's behavior in between the first and second period. Participants now have knowledge of actual peer behavior (as opposed to merely a prior belief) that is either pro- or anti-social. Consequently, individuals now face a trade-off decision in which deviation from the individual inherent characteristic θ^i has to be weighed against deviating from the observed peer behavior x_j^i. Because no peer behavior

information was given in period 1, individuals maximize:

$$\frac{\partial U_t^i}{\partial x_t^i} = 2\alpha_t^{ij}(x_1^j - x_t^i) - 2(\alpha_t^{ij} - 1)(\theta^i - x_t^i) \Leftrightarrow x^* = x_t^i = \theta^i - \alpha_t^{ij}\theta^i + \alpha_t^{ij}x_1^j \tag{4}$$

yielding the comparative static:

$$\frac{\partial x_t^i}{\partial \alpha_t^{ij}} = x_1^j - \theta^i \tag{5}$$

We can infer that a change of x_t^i in t depends on the social proximity as follows:

$$\frac{\partial x_t^i}{\partial \alpha_t^{ij}} = \begin{cases} > 0, & \text{if } \Psi_t^{ij} > \theta^i \ (x_t^i \text{ increases in } \alpha_t^{ij}) \\ = 0, & \text{if } \Psi_t^{ij} = \theta^i \ (x_t^i \text{ unaffected by } \alpha_t^{ij}) \\ < 0, & \text{if } \Psi_t^{ij} < \theta^i \ (x_t^i \text{ decreases in } \alpha_t^{ij}) \end{cases} \tag{6}$$

In order to account for Manski's (1993) reflection problem in studying the magnitude of behavioral contagion in stage 2, one has to hold individual j's behavior from period 1 constant while giving individual i the ability to revise her initial decision. The resulting difference in a peer's initial and individual i's revised behavior will be referred to as an contagion gap. From the previous results, we can easily infer that the larger (smaller) the α, the closer (further away) x_2^i is to (from) x_1^j and hence the smaller (larger) the gap.

Proposition: *Equation (4) provides a solution to the maximization problem and reduces the contagion gap to:*

$$|x_2^i - x_1^j| = |\theta^i - \alpha_2^{ij}\theta^i + \alpha_2^{ij}x_1^j - x_1^j| = |\theta^i(1 - \alpha_2^{ij}) + x_1^j(\alpha_2^{ij} - 1)| \tag{7}$$

Observe that, unlike in Akerlof's (1997) general conformity model, this approach generates a unique equilibrium prediction due to restrictions put on the social proximity parameter in the linear reaction function of the form $0 < \alpha_t^{ij} < 1$.

Section B: Additional Figures and Robustness Checks

Table A1: Descriptive Statistics

	Treatments		
	Baseline (Unknown Proximity)	Treatment 1 (High Proximity)	Treatment 2 (Low Proximity)
Participants	83	68	76
Female (%)	60% (0.49)	54% (0.5)	59% (0.5)
Age (average)	24.0 (4.86)	23.9 (3.76)	23.2 (2.80)
Average Amount (%) of Pro-Social Behavior (Before Observation)	50.1% (32.8)	47.9% (33.1)	51.4% (30.7)
Average Amount (%) of Anti-Social Behavior (Before Observation)	-66.3% (32.8)	-63.9% (32.8)	-57.1% (32.8)

Notes: Standard deviations in parentheses. Across all treatments, there are no significant differences among all characteristics that can be compared prior to the treatment manipulation. Average amount (%) refers to the amount given to (pro-social) or taken away (anti-social) from the charity relative to the maximum amount that a participant was able to give and take away, respectively.

The histogram in Figure A1 indicates that prior to observing their peer's behavior, active participants' decisions mainly clustered around 0%, which represents the decision to not change the initial equal distribution between oneself and the charity. After the observation, however, we notice a perceptible skew towards anti-social behavior, particularly in the high proximity condition. This finding provides a first indication that anti-social behavior is more contagious than pro-social behavior.

An alternative analysis allows us to break down exposure to peer behavior relative to the individual's initial behavior.

- **Observed Better Behavior:** Contains situations in which the active participant observed behavior of a passive participant that was better than his/her own initial behavior. This includes situations in which the observed behavior was either less harmful or more beneficial to the charity than one's own behavior.
- **Observed Worse Behavior**: Contains situations in which the active participant observed behavior of a passive participant that was worse than his/her own initial behavior. This includes situations in which the observed behavior was either more harmful or less beneficial to the charity than one's own behavior.

As indicated in the figures below, all of our previous results remain robust to this alternative analysis.

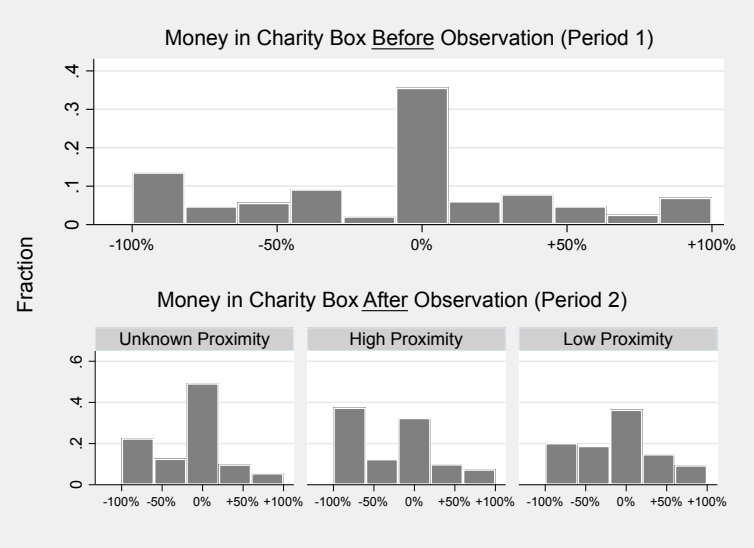

Figure A1: Incidence of choices to not change the initial equal distribution / give money to / take money from the charity before and after peer observation. The horizontal axis depicts a continuum of money left in the respective charity's cash account, with 0% representing the status-quo allocation. The vertical axis depicts the fraction of participants engaging in the particular behavior.

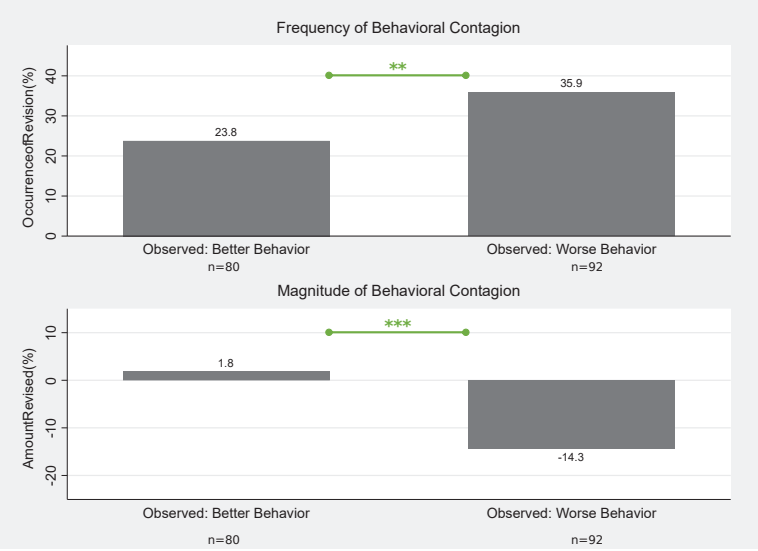

Figure A2:Robustness check of Figure 3. Frequency and magnitude of behavioral contagion by observation of pro- and anti-social behavior.

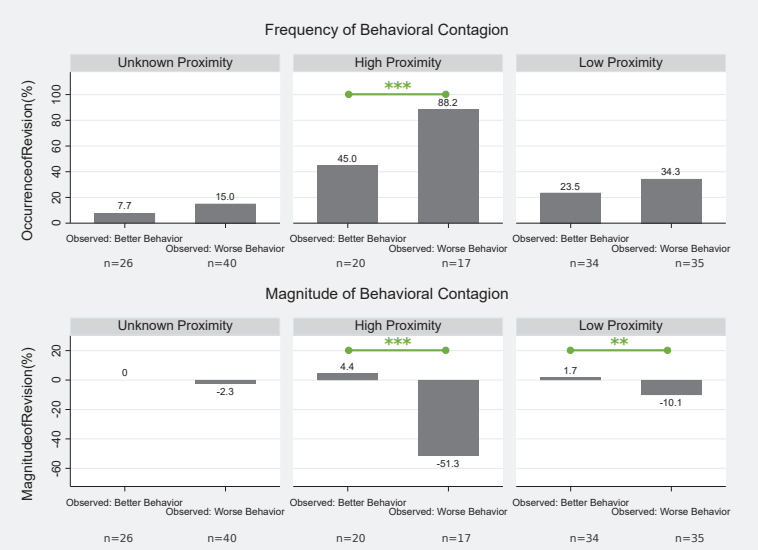

Figure A3: Robustness check of Figure 5. Frequency and magnitude of behavioral contagion by social proximity.

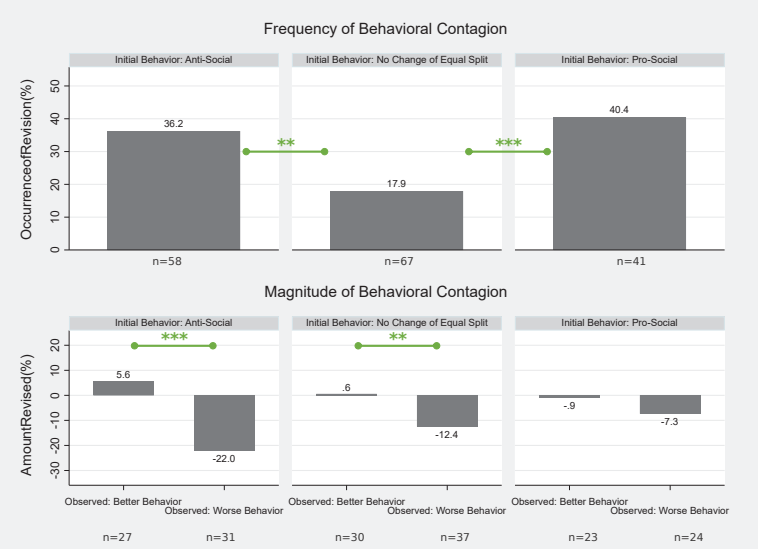

Figure A4: Robustness check of Figure 6. Frequency and magnitude of behavioral contagion by initial behavior, observed better/worse behavior and social proximity.

Section C: Additional Material

Instructions

General Information on the Experiment

- First of all, we would like to thank you very much for participating in this experiment. Please read the instructions carefully. The experiment will last for about 45-60 minutes.
- During the entire experiment, no communication is allowed. If there is something you do not understand or if you have any questions, now or at some point during the experiment, please raise your hand and remain seated. One of our colleagues will come to you and answer your question.
- During the experiment, you have the possibility to earn money. The amount you will receive at the end of the session depends on how many "Experimental Currency Units" (ECU) you earn during the experiment.
- At the end of the experiment, the amount of "ECU" that you have earned will be converted into real money at an exchange rate of 20 ECU = 1 Euro.
- All decisions you make during this experiment will remain anonymous. None of the participants gets to know the identity of other participants in the experiment and decisions cannot be linked to a specific participant. Moreover, you will be paid anonymously at the end of the experiment.

Order of Events:

- The experiment consists of a list of statements that you will receive at the beginning and further decisions. Explanations and information related to these decisions will be given as the experiment progresses. You will make these decisions once.
- Both you and the charitable organization of your choice (i.e. an officially registered charity organization) will be provisionally assigned a monetary amount of 300 ECU each.
- During the experiment you will have to decide on whether you want to:
    1. take a part or all of the money from the charitable organization.

2. leave the division of the sum of money as it is.
3. give a part or all of your money to the charitable organization.

- In case you decide to take money from the charitable organization, the respective amount of money will be transferred to your individual cash account and exactly the same amount will be deducted from the cash account of the charitable organization.
- Should you decide to give money to the charitable organization of your choice, the respective amount of money will be deducted from your individual cash account and given to the charity. The experimenter will double all ECUs remaining in the charity's account at the end of the experiment.
- Your decision remains anonymous and neither the other participants of the experiment nor the experimenters have the possibility to assign your choices to your identity.
- At the end of the experiment, one participant will be chosen at random and his or her choice will be implemented and count towards the charity (i.e. that choice will be relevant for the payment). The receipt of this donation will be published on the homepage of the BaER-Lab (http://www.baer-lab.org) in a timely manner. All other participants will receive 150 ECU (including the show-up fee) at the end of the experiment.

The total payoff of the participants:

- **In case you are the randomly chosen participant**
  300 ECU +/- (the amount of money that has been given to/taken from the cash account of the charitable organization)
- **In case you are not the randomly chosen participant**
  150 ECU

The total payoff of the charitable organization:
Amount of money in the cash account of the charitable organization of the randomly chosen participant

- At the end of the experiment, the relevant information on the payment will be made visible to each participant on his or her screen.
- After the actual experiment concludes, we will ask you to fill out a questionnaire. Please fill out the questionnaire carefully and truthfully.

Screenshots of Decision Screens

1. **List of statements: generates the proximity measure in all treatments**

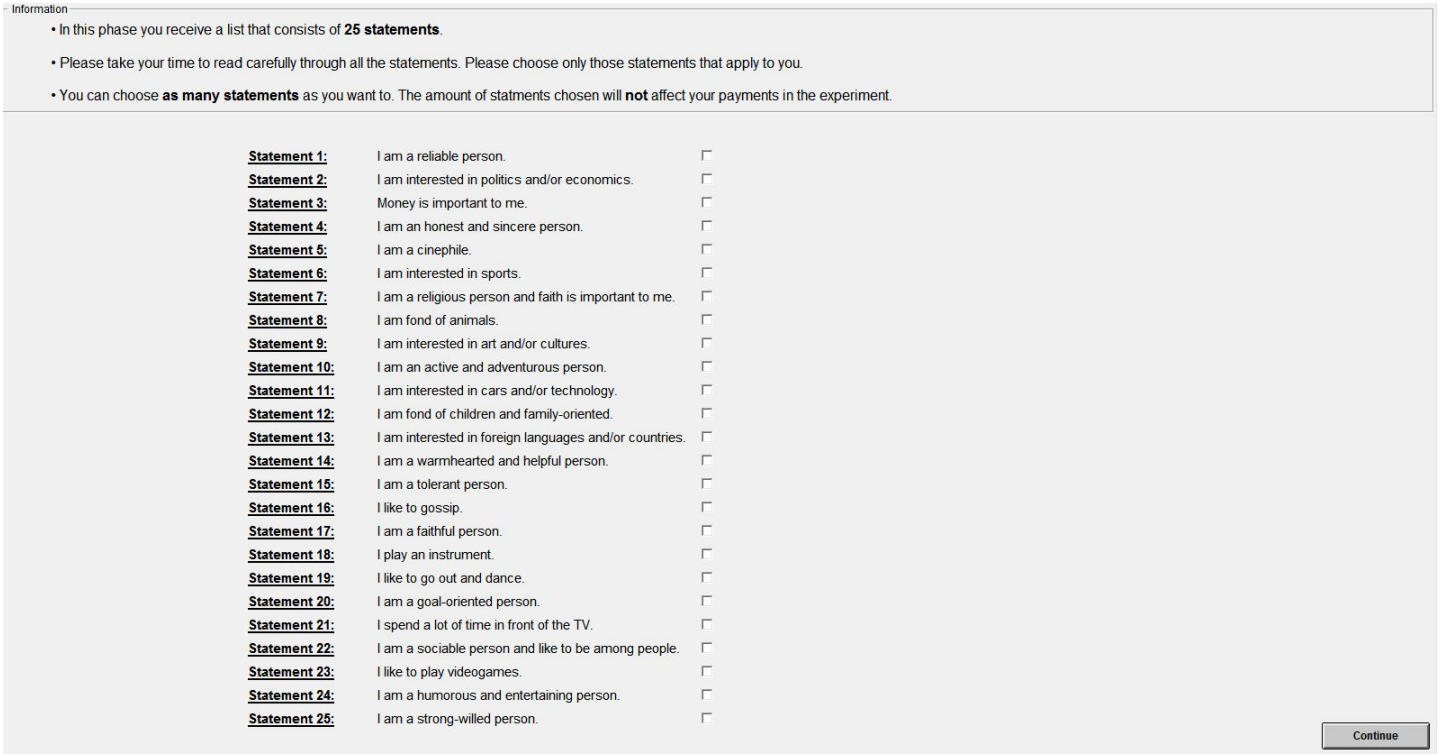

2. **First decision towards charity**

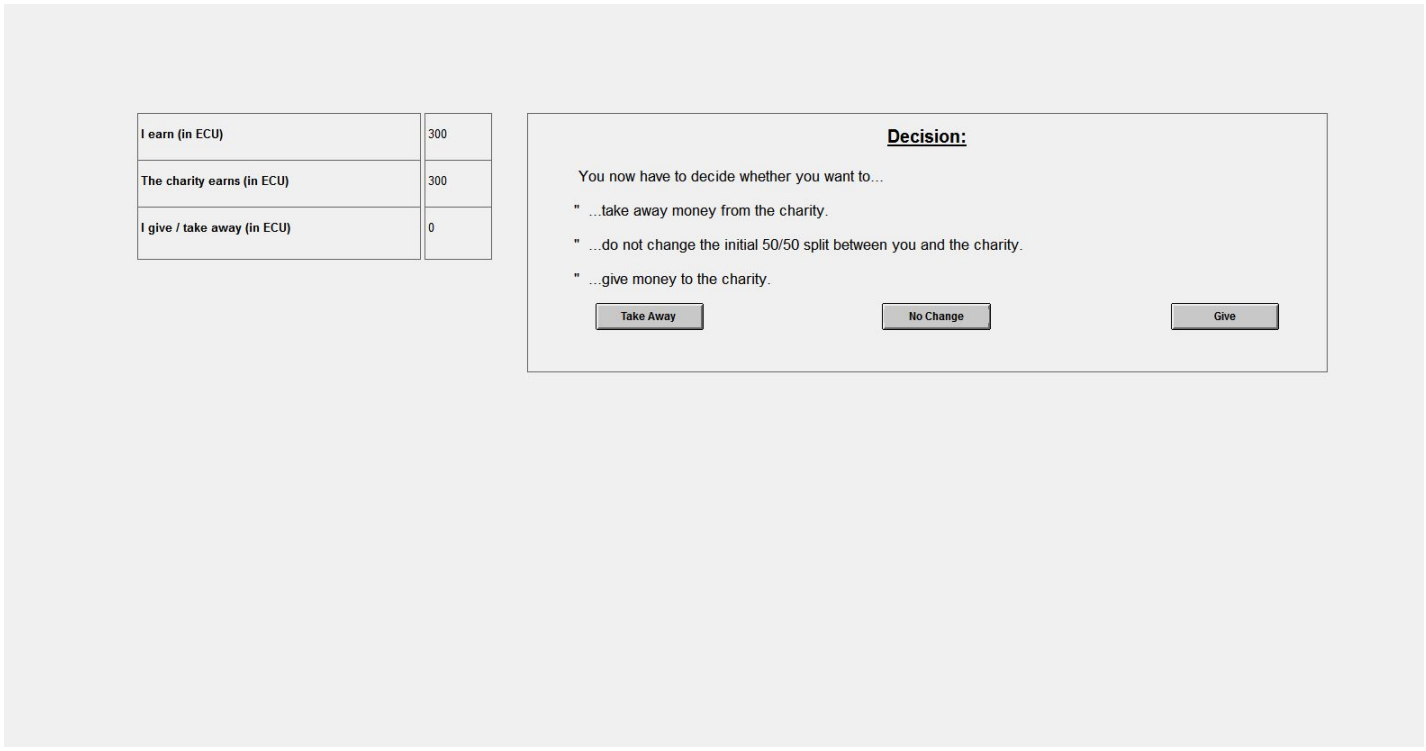

(Exemplary for the taking away decision)

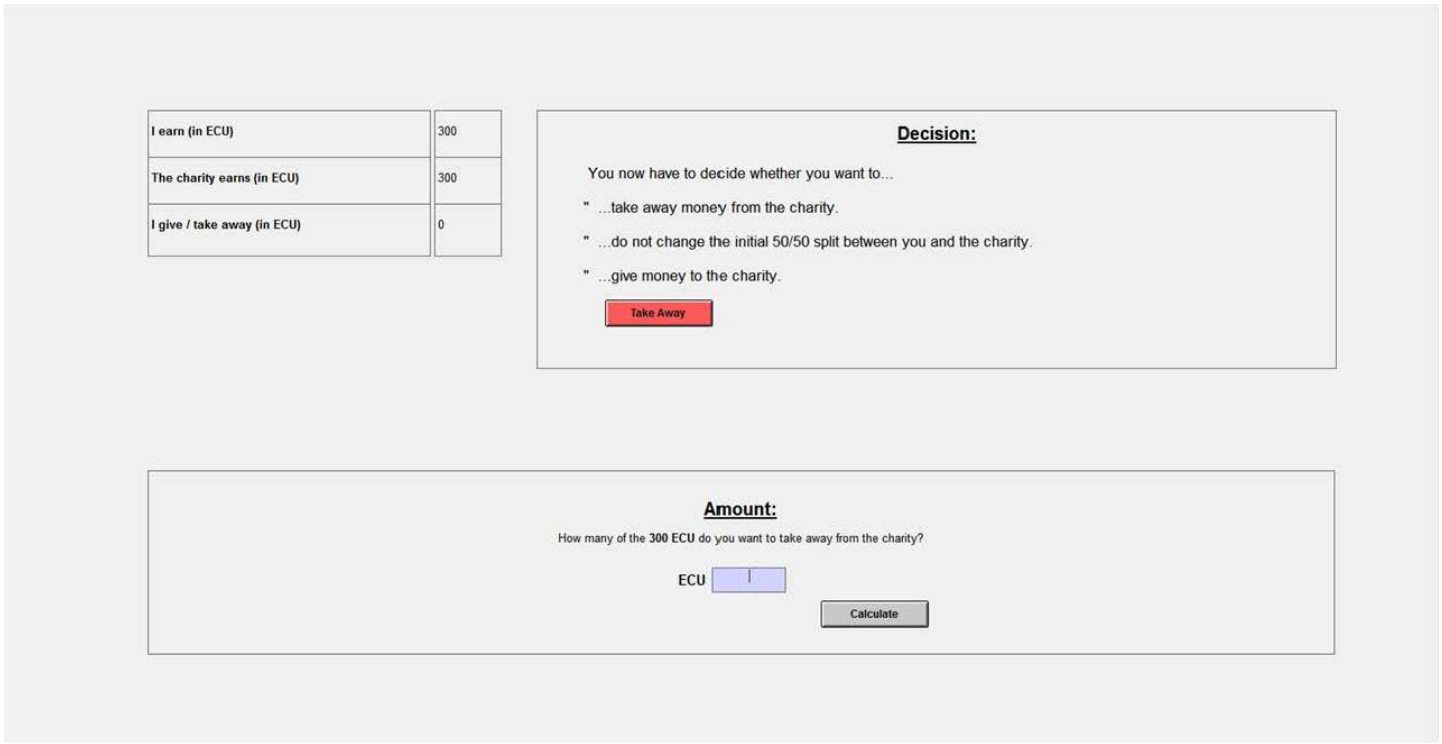

3. **Observation of one's peer and potential revision of one's initial decision**
   Important: treatments vary by the information on the social proximity measure (first box): unknown similarity (Baseline), more similar (Treatment 1), less similar (Treatment 2).

(Exemplary for the taking away decision)