CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The University of Nottingham
UNITED KINGDOM · CHINA · MALAYSIA

Discussion Paper No. 2017-06

Eugen Dimant

April 2017

On Peer Effects: Contagion of Pro- and Anti-Social Behavior and the Role of Social Cohesion

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
Fax: +44 (0) 115 95 14159
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

On Peer Effects: Contagion of Pro- and Anti-Social Behavior and the Role of Social Cohesion

Eugen Dimant*

University of Pennsylvania & Harvard University

This Version: April 2017

ABSTRACT:

Little is known about the underlying mechanisms of behavioral contagion, in particular with respect to differences in contagion of pro- versus anti-social behavior. Our principal contribution is the use of a novel experimental approach that enables us to analyze the contagion of behavior under varied levels of social distance to peers and differences in contagion of pro- and anti-social behavior. Anti-social behavior is found to be more contagious and social distance particularly drives the contagion of anti-social but not pro-social behavior. The results yield policy implications with regards to designing effective nudges and interventions to facilitate (reduce) pro- (anti-)social behavior, in both social and work environments.

JEL: C91; D03; D64; D81

KEYWORDS: Anti-Social Behavior, Behavioral Contagion, Charitable Giving, Peer Effects

---

*E-mail: edimant@sas.upenn.edu. This work has greatly benefited from conversations with Edward Glaeser, Daniel Houser, Lawrence 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, Elena Katok, Judd Kessler, and Robert Kurzban during my (visiting) research positions at the Harvard University, the University of Pennsylvania, and the University of Texas at Dallas, respectively. I also want to thank James Andreoni, Dan Ariely, Gary Charness, Robert Cialdini, Rachel Croson, René Fahr, Uri Gneezy, Burkhard Hehenkamp, Karla Hoff, Ilyana Kuziemko, John List, Rosemarie Nagel, Imran Rasul, Arno Riedl, and Tim Salmon as well as the participants at the London Experimental Workshop, the Norms, Actions and Games Workshop at the Institute for Advanced Study in Toulouse, the CESifo Venice Summer Institute 2016 on the Economics of Philanthropy, the Norms and Behavioral Change Workshop at Nottingham University, the Behavioral Science Hub Seminar at the London School of Economics, as well as numerous other workshops and conferences at UC Berkeley, UC San Diego, and Wharton, among others, for interesting discussions and helpful suggestions. Financial support by the German Research Foundation (DFG) through the SFB 901 is gratefully acknowledged.

1. Introduction

Individuals do not co-exist in pure isolation but interact within social contexts, or as emphasized by Elliot Aronson (2011): individuals are *social animals*. Despite the long tradition in anthropology and sociology, the study of the impact of norms, values, and social influence of peers is still a rather young field in economics. Fortunately, over the last two decades there has been a push in the economic discipline to expand our understanding of what comprises a more sophisticated individual decision maker by accounting for the individual's identity, morals, and other-regarding concerns (cf. Akerlof & Kranton (2000), Charness & Rabin (2002), Benjamin, Choi & Strickland (2010), Bénabou & Tirole (2011)). These approaches have enriched our understanding by considering social and economic decisions as a function of their respective social and economic environments as well as the relevance of peer behavior. Scholars in economics and psychology have attempted to shed light on the general mechanism of peer effects using controlled experiments. In the realm of theory (Baccara & Yariv, 2013) as well as in the lab and field experiments, peers are found to affect academic gains (Lavy & Schlosser (2011), Duflo et al. (2011)), investment decisions (Bursztyn, et al., 2014), littering behavior (Cialdini, et al., 1990), education (Sacerdote, 2001), adaptation of farming technologies (Bandiera & Rasul, 2006), productivity at work (Ichino & Maggi (2000), Falk & Ichino (2006), Mas & Moretti (2009), Azmat & Ichiberri (2010)), and juvenile behavior (Damm & Dustmann, 2014). In addition, the seminal Moving-to-Opportunity literature has broadened our understanding of neighborhood effects (cf. Case & Katz (1991), Katz et al. (2001), Kling et al. (2005)).

Most existing studies have focused on analyzing either pro-social or anti-social behavior in isolation (see discussion in Section 2). Our study extends the existing literature on peer effects in a number of ways. In addition to providing a clean approach to test *whether* contagion through peers exists in the context of both anti-social and pro-social behavior, we also deliver sound evidence on *how* such peer effects depend on the anti-/pro-sociality of observed peer behavior and the social coherence with one's peers. From a content perspective, our work focuses on better understanding the drivers of behavioral contagion in the domain of both pro- and anti-social behavior simultaneously. In particular, we analyze how social proximity to

one's peers mediates behavioral contagion. From a methodological perspective, our contribution is to introduce an approach that allows us to saliently induce and incrementally measure differences in social proximity and its impact on behavioral contagion.

Our paper addresses three main questions: first, does contagion[1] of pro-social behavior differ from contagion of anti-social behavior? Second, what is the role of social cohesion with one's peers in facilitating the spread of behavior in either direction? Third, whom does contagion affect and how? We examine these questions along the dimensions of frequencies (i.e. how often does behavior spread over) and magnitudes (to what extent does behavior spread over) of contagion. Understanding the drivers of pro- and anti-social behavior is pivotal from a policy perspective to which we will return in the last section. To the best of our knowledge, our paper is the first controlled experimental study to examine peer effects in pro- and anti-social behavior simultaneously. In addition, we provide a novel way to study such peer effects under varied levels of social proximity in a controlled setting, which allows us to enrich the scope of this contribution.

Embedded in a controlled lab experiment, our approach to measuring peer effects is a variation and extension of the dictator game as introduced by List (2007) and Bardsley (2008). In our sequential setting, participants are given the opportunity to give or take money away from the charity, first in isolation and then after observing peer behavior, which is introduced in the form of a revision option. Revising behavior into the direction of observed peer behavior is accounted for as *behavioral contagion*. In more detail, our design follows a straightforward procedure: action – peer observation – reaction. In this variant of a two-player dictator game, the participant (dictator) is matched with a charity (recipient). The participant's action space is characterized by: (i) taking money away from the charity, (ii) retain the status-quo of an equal distribution, or (iii) giving money to the charity. After this initial decision, participants are randomly divided into active and passive participants, with the difference being that while the former get to observe the initial behavior of the latter and decide whether to revise their own initial decision, the behavior of passive players is held fixed

[1] Behavioral contagion describes the spillover of behavior through social interaction. More precisely, it is 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).

and they are not given a revision option. This design choice allows us to account for obstacles that are typical for peer effect studies, such as endogeneity and reflection (Manski (1993), Manski (2000), Angrist (2014)). Alongside the observation of a peer's behavior, treatment variations include the alteration of unveiled social proximity information to the peer based on the similarity score calculated from overlapping answers to the personality questionnaire.[2]

In our experiment, social cohesion is derived from preference similarity for things such as food, sports, or family, among other things. Arguably, common preferences and/or interests are the first step in creating a common identity. For this purpose, we introduce an approach that allows us to induce, measure, and gradually vary the proximity between peers in order to study behavioral contagion in different social settings. We construct a list of personal statements taken from a major American dating website to categorize participants according to overlaps in preferences and interests. We then use that array of responses to create a proximity index (matching score) that indicates the number of answers that the respective participants have in common and use this score to vary the proximity signal observed by the participant. Such a measure of social cohesion can be used as an exogenously varying matching device to study peer effects in the lab, which has not been possible with previous proximity measures. This approach creates three major advantages: for one, the controlled laboratory setting mitigates the previously discussed arising difficulties when studying peer effects in the field.[3] For another, dating website questions are found to represent reliable matching approach (for a discussion, see Hitsch, Hortaçsu & Ariely (2010)), thus validating our approach. Along similar lines, such a measure of social proximity allows us to mimic social affection that is normally absent in the artificial laboratory setting due to its intended nature of anonymity. Lastly, this

[2] In economics, the value of social information has previously been studied by Bohnet & Frey (1999), Charness, Rigotti & Rustichini (2007), Charness & Gneezy (2008), Chen & Li (2009), Benjamin, Choi & Strickland (2010), Eckel & Petrie among others. For a review see Costa-Font & Cowell (2015)). Although not strictly in line with our approach, the general idea of measuring peer effects using varying amounts of observability to overcome the reflection problem is inspired by the work of Bandiera et al. (2005) and Mas & Moretti (2009).

[3] In a recent meta-study, Herbst & Mas (2015) found that laboratory studies on peer effects generalize quantitatively to findings in naturally occurring environments in the field, rendering both approaches to be important complements in the study of behavioral spillovers.

approach is suited to vary social proximity among individuals gradually beyond a binary distinction of proximity, which is not easily accomplished with the standard minimal group paradigm approaches.

Briefly summarizing our results, we find that behavioral contagion is driven by both exposure to behavior and the social cohesion to the peer, with behavioral contagion being asymmetrically biased towards the spread of anti-social behavior. Overall, our results suggest that within a given peer context it is more likely to observe behavioral contagion in the form of anti-social than pro-social behavior. Across different specifications, we find that social proximity to the peers is more relevant to the crowding-out than to the crowding-in of pro-social behavior. The interaction between social proximity and type of observed behavior adds to the understanding of peer effects and yields relevant policy implications, which we will discuss in more detail in the concluding section.

The paper is structured as follows: we discuss relevant literature in Section 2, present the experimental design alongside with the hypotheses in Section 3, the results in Section 4, and discuss the implications in Section 5. We conclude in Section 6.

2. Literature Overview

Several economic and psychological theories speak to reasons why individuals comply with peer behavior. The economic concepts include *social decisions and social distance* (Akerlof (1997), Glaeser & Scheinkman (2004)), *image related concerns and taste for conformity* (Bernheim, 1994), and *imitation of behavior or preferences* (Alós-Ferrer & Schlag (2009), Sliwka (2007)). Some of the (social) psychological concepts are *social learning* (Bandura, 1971) and *norms* (Cialdini et al. (1990), Bicchieri (2006)). Other seminal research indicates that peer pressure exhibits shame and guilt, which can translate into behavioral contagion (Kandel & Lazear, 1992), social pressure (DellaVigna, et al., 2012), and identity priming (Kessler & Milkman, forthcoming).[4]

[4] Many of these concepts are not strictly distinct, in both their assumptions and predictions. In this paper, we will not attempt to resolve which approach explains behavioral contagion best but rather focus on shedding light on the drivers of behavioral contagion and its interrelation with the social identity dimension. See Dimant (2015) for a discussion of these channels and their implications on behavioral contagion.

Economic research has highlighted behavioral spillovers in pro-social behavior that extends to voluntary contribution to a public good, cooperation, and reciprocity (Andreoni (1995), Offerman (2002), Frey & Meier (2004), Croson & Shang (2008), Shang & Croson (2008), Bicchieri & Xiao (2009), Chen et al. (2010), Gächter et al., (2013), Thöni & Gächter (2015)), and within the unethical and anti-social domain, such as doping (Gould & Kaplan, 2011), dishonesty (Gino, Ayal & Ariely (2009), Innes & Mitra (2013)), and stealing behavior (Falk & Fischbacher, 2002).[5] We discuss the literature closest to our experiment shortly (see Dimant (2015) for an extensive overview).

In economics, conclusions of limited experimental research have also pointed to the contagion of both selfish and dishonest behavior. Bicchieri and Xiao (2009) studied the effect of a dictator game with varying information on other participant's selfish or fair behavior, finding that fairness in actions is contagious. In a more delinquent context, Falk & Fischbacher (2002) investigate peer effects in the form of conditional stealing behavior. Their findings suggest that, on the aggregate level, people make stealing decisions conditionally based on the behavior of their peers. Innes & Mitra (2013) use a variant of Gneezy's (2005) deception game to study whether dishonesty breeds dishonesty. Their findings suggest that already the beliefs about other's dishonesty are contagious, driven by the wiggle-room created by such social cues and creating a justification device for one's personal dishonest behavior.

Several field experiments have investigated the change of individual contribution levels in response to the observation of other people's contribution decisions. Exemplarily, Cialdini et al. (1990) show that the observation of another person's behavior leads to less littering. What is more, Offerman (2002), Shang & Croson (2008), and Croson & Shang (2008) assess the effect of downward and upward social information in contribution decisions to fund public goods. Their results indicate that the downward adjustment of pro-social behavior (own contributions to the public good) is larger than the upward adjustment. While these studies focus on the extent of pro-social behavior in isolation, our study capitalizes on a setting that allows us to study the extent of pro- and anti-social behavior simultaneously.

[5] The contagion of behavior through social contact has also been shown in other disciplines, such as network theory (cf. Centola & Macy (2007)).

A seminal line of controlled experiment-in-the-field research has provided robust evidence on the effects of neighborhoods and peers on the development of children, criminal attitudes, and education, among other things. Case & Katz (1991) found evidence for criminal behavioral contagion within both families and neighborhoods in the Boston area. Subsequently, multiple research projects have examined the short- and long-run effects of the Moving-to-Opportunity (MTO) project in which families are eligible to participate in a lottery for vouchers that would potentially help them to move to a better and safer neighborhood. Katz, Kling & Liebman (2001) examined the short-run effects of the MTO project on the well-being of the families who were offered a voucher. Their findings indicate a substantial improvement of well-being along different dimensions, including increased safety, and improved health conditions both mentally and physically. Surprisingly, young men were especially susceptible to the neighborhood change, while the young women's disobedience remained invariant. For criminal behavior, Kling, Ludwig & Katz (2005) found a strong gender effect. In terms of reduced arrests for violent crimes, men react positively to improved living conditions in the short-run. In the long-run, however, these effects vanish. Contrary to expectations of moving to a better neighborhood, males' general problem behavior and property crime arrest soar irrespectively. Conversely, females' criminal behavior decreases. Capitalizing on a different but comprehensive dataset that includes the assignment of refugee immigrants to Denmark from 1986 to 1998, Damm & Dustmann (2014) find that the share of convicted young people in the neighborhood significantly increases both the probability for a male's convictions later in life and the total number of convicted crimes executed by men. Their findings suggest that the spillover-effects of neighborhood crime are distinctively linked through the channel of social interaction, which is, however, only true for youth criminal behavior. All in all, these studies provide robust evidence not only for peer effects in general but particularly within the domain of criminal behavior.

3. The Experiment

3.1 Experimental Design and Procedure

Existing approaches face various challenges in studying peer effects in a controlled setting (Manski (1993), Manski (2000), Angrist (2014)). We deal with the noted reflection problem

by introducing a novel design, which we will discuss in more detail in the design section of our paper. In short, our approach centers on two key design elements: firstly, only those who *actively observe* another participant's behavior can react and revise initial behavior.[6] Secondly, behavior of those who *are observed* is held fixed and cannot be altered after the initial decision (public knowledge), thus using a lower-bound approach to induce different levels of social proximity. The additional information that an observing participant in the high (low) proximity treatment receives alongside observing the peer's actual behavior is that the preference similarity to this observed peer, as calculated based on overlapping answers to the personality questionnaire, is higher (lower) than to the second (unobserved) peer. Such an approach allows us not only to study general peer effects in an unbiased way, but also to shed light on the relevance of factors such as the social proximity to one's peers and the pro-/anti-sociality of observed behavior in driving behavioral changes. This experimental set-up allows us to contribute to existing research on peer effects and opens a venue for future research.

We mimic social proximity by using statements taken from a major American dating website that capture individual preferences and interests. We use the matching scores of overlapping answers among lab participants as an exogenous matching device across treatments. This allows us to study decision-making beyond simple ingroup - outgroup comparisons. Rather, our approach provides us with an extensive array of possibilities to match participants according to their shared similarities. To the best of our knowledge, we are the first to use such an approach. Thus, we not only complement existing field studies, but also broaden the scope and utilization of lab experiments in explaining behavior and behavioral changes in peer settings, especially within the anti-social domain.

Basic Procedure

Consider a variant of a one-shot dictator game in which the participant (dictator) is matched with a charity (recipient). The dictator's action space entails (i) taking away money from the

[6] In our study, we report lower-bound peer effects because adaptive behavior remains unobservable by one's peers and thus carries no signaling value. Because such a setting distinguishes us from what is typically meant by the term conformity, we use the more encompassing term behavioral adaptation. For a discussion of the mechanism and empirical literature on conformity, see Bernheim (1994) and Bikhchandani, Hirshleifer & Welch (1998). See Wheeler (1966) for a discussion and differentiation from other frequently used terms such as conformity and imitation.

charity, (ii) retaining the distributional status-quo, or (iii) giving money to a charity. In following Eckel & Grossman (1996), we use a charity to increase the saliency of decisions.[7]

The experiment is played one-shot with a possibility to revise one's initial behavior. The implemented design is a straightforward single iteration of an initial stage 0 at which participants filled out a personality questionnaire and three stages: **1st stage** - action (initial decision) → **2nd stage** - peer observation → **3rd stage** - reaction (revision decision). In between the initial decision and potential revision, individuals can observe the initial behavior of another random participant. Each experimental session followed all stages (0 - 3). Subjects were randomly assigned to treatments within each session. In all sessions, the dependent variable is the final choice of giving to/taking from the charity. This study tests whether the revision choices change when observing another dictator choice conditional on both the behavior of and the proximity to the observed dictator.

Prior to First Decision:

Participants answer a personality questionnaire containing 25 items that were taken from a major US dating website (see Appendix for details). The answers to these questions will then be used to calculate a proximity score among participants in the second stage. Depending on the treatment, the participants may or may not observe the matching score in addition to the dictator decision of the peer.

1st Stage - The Action:

Starting with an equal distribution of money, each individual decides whether to (i) donate own money to the charity's account, (ii) maintain the initial equal distribution, or (iii) take money from the charity and keep it for themselves. Note that in order to exclude any hedging concerns throughout the experiment, information about the specifics of the design were only provided where necessary in order to reach a deliberate decision. That is, at Stage 1 participants were neither aware of the possibility to observe peers nor to revise their initial decision, ensuring unbiased initial behavior.

[7] In order to reduce biases and increased wiggle-rooms for misbehavior stemming from forcing a particular charity on the participants, instead 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.

2nd Stage – The Observation:

After all participants made the initial dictator decision, exactly two participants were chosen from the pool and assigned the role of the *passive players*. The remaining n-2 participants were assigned the role of the *active players*.[8] At this stage, each active player is given the chance to observe one passive player at random, who has engaged in either pro-social or anti-social behavior. In all three treatments, active players can always observe the initial 1st stage dictator behavior of exactly one passive player. Treatment differences lie in whether or not the active participants are observing additional information on the social proximity score with the passive player as calculated based on the amount of overlapping answers to the 25 item personality questionnaire. Except in the Baseline, in which proximity information remained unobserved, observers (active players) received additional information on the social proximity to the observed passive peer. In the proximity treatments, the observation of a passive participant includes information on whether the proximity to the observed passive player is higher or lower compared to the other (unobserved) passive player. That is, at no point does the active player receive information about the *exact* matching score with the observed passive player. Instead, the active player receives information about the *relative* matching score, indicating that the matching score with the observed passive player is higher or lower than with the second (but unobserved) passive player.[9]

[8] Importantly, the active peers received the information that the passive players will always consist of one participant who has money given to and one participant who has taken money from the charity, but were not told any additional information regarding the population of behavior in the session. The randomly chosen passive participants were chosen from the extremes of both ends (those who gave everything to and took everything from the charity), of which there were always multiple participants in each session. We resort to this approach to avoid biases stemming from social learning or updated beliefs about distribution of anti-/pro-social behavior in the population. A random participant drawn each from the anti-/pro-social spectrum contains no signaling value on the true fraction of anti-/pro-social participants in the population, thus retaining the salience of the observation. 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.

[9] The implementation of the low- and high-proximity information followed a very straightforward calculation. For each participant of the active group, an individual proximity score to both participants of the passive group was calculated based on overlapping answers in the list of statements. From each active participant's individual perspective, the passive participant with the higher (lower) score was labeled as the high (low) proximity peer. In fact, this calculation approach allows for the same passive person to be of high (low) proximity to one active person, while being of low (high) proximity to another active person, thus truly randomizing information. We thus report lower-bound results and abstained from providing explicit matching scores or percentages to retain maximum control. In addition, this allows us to take care the *false-consensus* effect, in which people systematically overestimate the degree of similarity to others. The provision of social cues of this kind allows the participants to update their beliefs reliably with respect to the actual degree of similarity (Ellingsen & Johannesson, 2008).

3rd Stage – The Reaction:

After observation, the active player is given the choice to revise his initial decision. Conditional on whether the active player chooses to revise, the revision decision will then replace the initial decision. The experiment ended after this stage, no other participant was able to observe the revision decision.

The experimental design is detailed in figure 1.

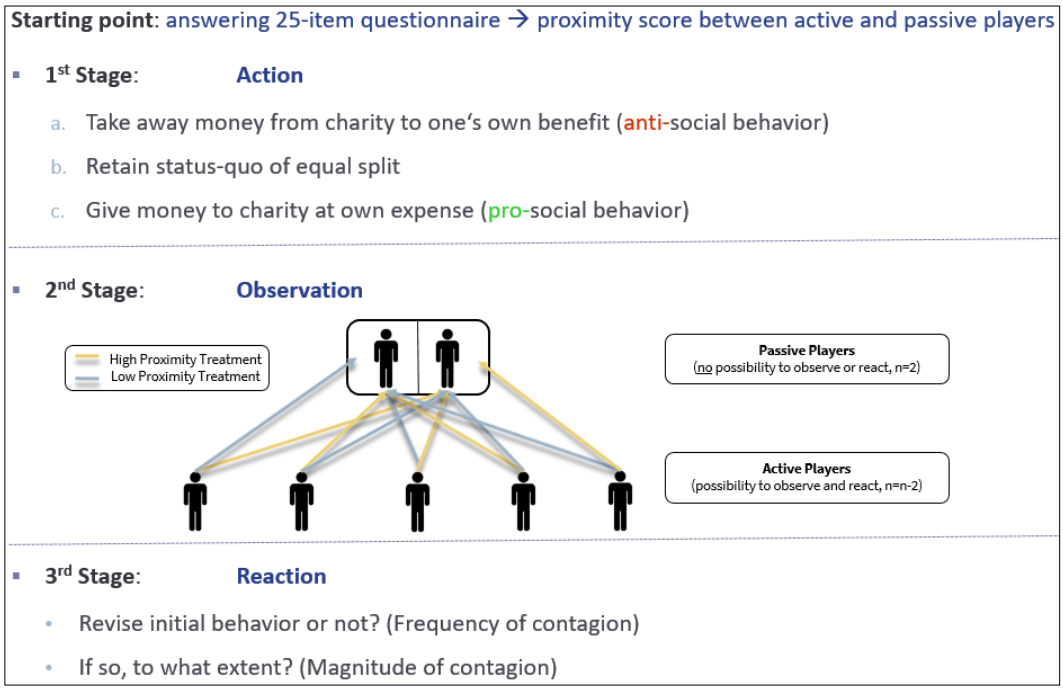

Figure 1: *Experimental Design and Procedure. In all treatments, the observed information is the passive player's initial dictator decision. Note that in the Baseline (unknown proximity condition), the player is unaware of the proximity score and hence does not know whether the observed passive player is of high or low proximity.*

The Design in More Detail

Alongside the actual behavior, treatment variations include the alteration of unveiled social proximity information of the observed participant. That is, in addition to learning actual behavior and the amount that was taken away from or given by this participant to the charity, information regarding the social proximity between matched participants is varied with the random treatment assignment. The treatment variation lies in the information given regarding the social proximity to the observed peers: (i) no information on proximity, (ii) high proximity,

and (iii) low proximity information. More specifically, the additional information that an observing participant in the high (low) proximity treatment receives alongside observing the peer's actual behavior is that the preference similarity to this observed peer, as calculated based on answers to the personality questionnaire, is higher (lower) than to the second (unobserved) peer.[10]

Proximity is calculated based on overlapping answers in the list of statements that participants filled out right before start of each session and then presented to the participants in the form of below- or above-average proximity information to the observed peer. We capitalize on a shortened 25 item list of statements compiled from a major US American dating website to ensure the validity of the questions in successfully matching people (see Hitsch, Hortaçsu & Ariely (2010)) for a discussion). The business concept of dating websites is based on achieving high matching success rates, thus using validated questions improves the success of ensuring social proximity between participants in the lab.

Several points are worth noting. In order to retain maximum control, in each session exactly two participants were chosen at random as *passive* (being observed by peers), while all other participants randomly chosen as being *active* (observing peers) always observe exactly *one* of these two passive players. The treatment differences are based solely on the social proximity information, i.e. whether the observing active participant has a higher matching score with the observed passive player as compared to the unobserved passive player. Participants knew that the observed social proximity to that peer will depend on the treatment one has been randomly assigned to and the number of overlapping answers to the initial personality questionnaire. The observer would then learn the behavior (in all conditions) and social proximity (if not in the Baseline condition) of exactly one passive player.

What is more, to improve comparability and reduce any potential confounds, each experimental session contained all three treatments. Participants were randomly assigned to one of the three treatments within the same experimental session. The experiment concluded with a battery of non-incentivized questions to elicit attitude towards risk and greed, among others.

[10] We provide screenshots detailing the exact wording in the Appendix.

Payoff structure:

Importantly, to exclude any form of strategic interaction that might potentially dilute results or affect their saliency, the participant's decisions only affected one's own and the chosen charity's payments but not those of other participants. That is, each individual's decision had no monetary impact on other individuals, making a change in behavior due purely to behavioral contagion and not to other-regarding concerns.[11] We make this even more salient by randomly picking one of the individuals at the end for which the behavioral decision was implemented, while everyone else received a flat income irrespective of his actual behavior. In monetary terms, each participant and the respective charity received the ECU equivalent of 15 Euro, thus allowing a participant to leave with a maximum (minimum) of 30 (0) Euro if the participant decided to take away all the money from (give all the money to) the charity. In order to increase the saliency of pro-social behavior, we added a multiplier to the setting. That is, the experimenter doubled all Euro remaining in the charity's account at the end of the experiment. At the end, one participant was chosen at random and the respective decision was implemented with respect to taking from or giving to the charity. The remaining participants in the same session received a flat income of 7.50 Euro.

3.2 Hypotheses

In order to generate hypotheses that align with existing theory, we derived our predictions from previous research.[12] Following our motivation, our hypotheses will focus on (i) whether behavioral contagion is symmetric on either side of the anti-social and pro-social spectrum, and (ii) on the mediating effect of social proximity to the observed peers.

So far, the economic literature has been fairly silent on answering two naturally arising questions with respect to peer effects and the spillover of behavior.

Question 1: *Is there a systematic difference in behavioral contagion with respect to anti-social behavior as compared to pro-social behavior?*

[11] To some degree, the experimental design embeds the considerations of Alós-Ferrer & Schlag (2009).

[12] The hypotheses can also be derived by extending Akerlof's (1997) model. See Dimant (2015).

It is reasonable to assume that behavioral contagion is a function of the nature of the observed behavior due to implied differences in costs. Good behavior implies bearing costs in order to improve the well-being of others, while bad behavior often implies improving one's own well-being in one way or the other at the expense of a third party. This reasoning is in line 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 & Shang (2008), Chen et al. (2010)) and the findings on slippery-slope effects (Gino & Bazerman, 2009). This assumption is also in line with the recent findings by Smerdon, Offerman & Gneezy (2016), Bicchieri, Dimant & Gächter (2017), and Nosenzo, Gächter & Molleman (2017) on the role of bad behavior and the perseverance of bad norms. In addition, research examining neighborhood effects indicates that the adaptation of behavior is asymmetric and depends on whether one has been exposed to good or bad influences and to what extent (Kling, et al., 2005). We thus derive our first hypothesis as follows:

H_1: ***Anti-social behavior is more contagious than pro-social behavior.***

Following the existing literature on social coherence, it is reasonable to assume that observing the behavior of people who are socially closer or similar depicts a more salient signal in terms of what is socially accepted or an existing norm (i.e., for the case of reciprocity, see Charness et al. (2007)). Pro- and anti-social behavior differs in terms of the information set available to the individual. While the former behavior might entail some uncertainty with respect to what is *appropriate* within a given context, the latter behavior is less ambiguous: the nature of anti-social behavior implies the overstepping of (social) boundaries or infringing upon laws. Here, the wiggle-room for self-justification is narrower. [13] This raises the question:

Question 2: *To what extent is behavioral contagion in either direction (i.e. pro- and anti-social behavior) mediated by the social proximity to the peers?*

Existing research indicates that social identification and proximity is a predictor of behavior in different contexts related to charitable giving, trust, punishment, and reciprocity (cf. Akerlof (1997), Charness et al. (2007), Chen & Li (2009), Leider et al. (2009)), as well as

[13] In his popular book, Dan Ariely (2013) shares a plethora of anecdotal evidence and supporting research studies that highlight the role of social factors affecting moral contagion and the transgression of unethical behavior.

neighborhood effects (cf. Damm & Dustmann (2014)). It has also been shown that norms or behavioral prescriptions are associated with one's identity, thus potentially rendering behavioral adjustments more likely in situations in which identity to peers is salient (Akerlof & Kranton (2000), Benjamin, Choi & Strickland (2010)). Results on advice seeking and giving substantiate the claim that norm signaling of what is socially acceptable is more salient if it comes from a person who is closer in the social dimension (Gino & Moore, 2007). Consequently, we derive our second hypothesis as follows:

H_2: ***Social proximity amplifies the contagion of behavior in general and anti-social behavior in particular.***

4. Results

We conducted the experiment at the BaER-Lab at the University of Paderborn, Germany. Participants were recruited using ORSEE (2004). We used zTree (Fischbacher, 2007) to run our experiment. In sum, 227 participants throughout 9 sessions were randomly assigned to one of the three treatments (unknown proximity, high proximity, and low proximity). Each session lasted about 45 minutes and the hourly average earnings were €10.50. An average of €30 was donated per session to various charities. We present descriptive statistics in Table 1.

	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)

Table 1: *Descriptive statistics. 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 (300 ECU) that a participant was able to give and take away, respectively.*

Next, we turn to testing our hypotheses by shedding light on the following two questions: Did behavior spread? And if so, to what extent? For this purpose, we will study behavioral contagion both in frequencies and in magnitudes. [14] In particular, the *frequency* of contagion is said to be affected (i.e. equals the value 1) if, after observing a peer's behavior, the participant revised his/her initial decision in the direction of the observed behavior. This behavior indicates the existence of behavioral contagion caused by peer effects. In contrast, for the *magnitude* of contagion, a value >0 indicates that the participant is becoming more pro-social (i.e. by reducing initial giving to or increasing initial taking from the charity), while a value <0 indicates that the participant is becoming more anti-social (i.e. by reducing initial taking from or increasing initial giving to the charity).

4.1 Effect of Observed Behavior

Our first approach to answering the two questions posed above is to look at the frequencies and magnitudes of behavioral contagion conditional on the type of observed behavior. Our results suggest that observing anti-social behavior is generally more contagious, both in frequencies and magnitudes of behavioral contagion. For this, we break down observed behavior in two ways:

i. **Observed: Better Behavior -** contains situations in which the observing (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.

ii. **Observed: Worse Behavior -** contains situations in which the observing (active) participant observed behavior of a passive participant that was *worse* than his/her own

[14] Importantly, in dealing with the reflection problem (Manski (1993) and Manski (2000)), participants are randomly assigned to passive and active players following their initial decision towards the charity. Out of the initial 227 participants, 18 (8%) participants were randomly assigned to being observed and thus remained passive after their initial decision. By definition, behavioral contagion can only be studied on the side of the active players; hence, our analyses will exclusively focus on them. In trying to reduce demand effects, our design choice closely follows Eckel & Petrie (2011). Here, 24 (11%) active participants decided to opt-out and not to observe peer behavior, suggesting that if there would have been demand effects, they would have likely been negligible. For a discussion on strategic ignorance, see Carrillo & Mariotti (2000), Dana, Weber & Kuang (2007), and Grossman (2014).

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.[15]

We first look at differences in frequencies of behavioral contagion, i.e. situations in which the active participant decided to revise his/her initial behavior in the direction of observed behavior, across these three different specifications (upper part of Figure 2). We employ a simple two-sample test on the equality of proportions and find that observing worse behavior triggers behavioral contagion more often than observing better behavior, with the differences being weakly significant (35.9% vs. 23.8%, p=0.08).

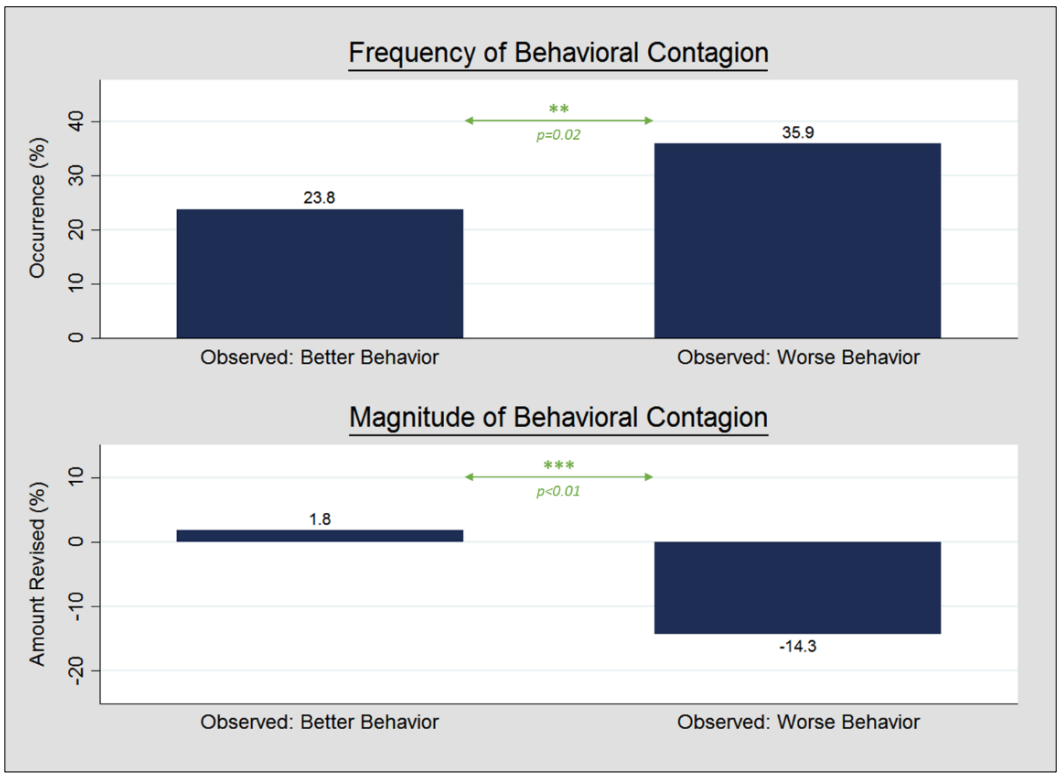

Figure 2: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on observed behavior. *Frequency of behavioral contagion* refers to the prevalence at which participants revised their initial behavior into the direction of observed behavior. *Magnitude of behavioral contagion* refers to the extent of the revised behavior relative to one's initial behavior. A positive (negative) number indicates that revised behavior was more pro-social (anti-social), i.e. participant gave more money to (took more money away from) the charity compared to one's initial behavior. Observed: better (worse) behavior indicates that the observed peer behavior was more pro-social (anti-social) than one's own behavior.

[15] Importantly, one alternative way of looking at the data is by breaking behavior down into two clear groups unconditional on one's own behavior: pro-social (when passive participant took a positive amount from the charity) and anti-social (when passive participant gave a positive amount to charity). A potential downside of such an approach is that one would lose the opportunity to observe heterogeneity in behavior when observations were more or less anti- or pro-social than one's own behavior. 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.

In a next step, we look at the magnitudes of behavioral contagion (lower part of Figure 2). We find that *after observing worse behavior*, participants become more anti-social by a rate of 14.3%. On the other hand, *after observing better behavior*, participants become slightly more pro-social by a rate of 1.8%. From this, it follows that with a ratio of almost 8:1 (= 14.3/1.8) the downward adjustment of behavior is more pronounced than the upward adjustment, indicating that anti-social behavior is more contagious than pro-social behavior. Using Mann-Whitney-U (MWU) statistics we find this difference to be significant at the 1% level (-14.3% vs. 1.8%, MWU, $p<0.01$).

Overall, our results are in strong support of hypothesis H_1 and indicate that behavioral contagion is indeed asymmetric, with anti-social behavior being more contagious than pro-social behavior. This is true for behavioral contagion in both frequencies and particularly in magnitudes.

4.2 Effect of Social Proximity

We parallel our previous analysis by shedding light on both frequencies and magnitudes in behavioral contagion conditional on the social proximity of the observed peer. As will be shown below, our results robustly indicate that higher social proximity indeed triggers stronger behavioral contagion, particularly contagion of anti-social behavior. When social proximity increases, the frequency and magnitude of revised behavior increases as well. In following our experimental design, we refer to the high proximity, low proximity, and unknown proximity conditions as follows:

i. **High Proximity -** the calculated matching score (based on the 25 item personal questionnaire) between the observing (active) participant and the observed (passive) participant is *higher* compared to the other unobserved (passive) participant. That is, in addition to observing the actual behavior of the passive participant, the active participant receives note that the matching score is higher compared to the matching score with the other passive but unobserved participant.

ii. **Low Proximity -** the calculated matching score (based on the 25 item personal questionnaire) between the observing (active) participant and the observed (passive) participant is *lower* compared to the other unobserved (passive) participant. That is, in addition to observing the actual behavior of the passive participant, the active participant receives note that the matching score is lower compared to the matching score with the other passive but unobserved participant.

iii. **Unknown Proximity –** the active participant randomly observes the actual behavior of one of the two passive participants but does not receive any information about the matching scores with either of the two passive participants.

We again first look at the *frequency* of behavioral contagion and observe that the contagion happens significantly more often in the high proximity condition, followed by the low proximity and unknown proximity condition. 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 two times as likely than in the low proximity condition (64.9% vs. 29.0%, $p<0.01$) and more than five times as likely than in the unknown proximity condition (64.9% vs. 12.1%, $p<0.01$).

What is more, 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, $p=0.02$).

From these results, we can conclude that high social proximity triggers significantly higher behavioral contagion not only in frequencies but also in magnitudes. This is particularly true for the contagion of anti-social behavior and thus supports our hypotheses H_1 and H_2.

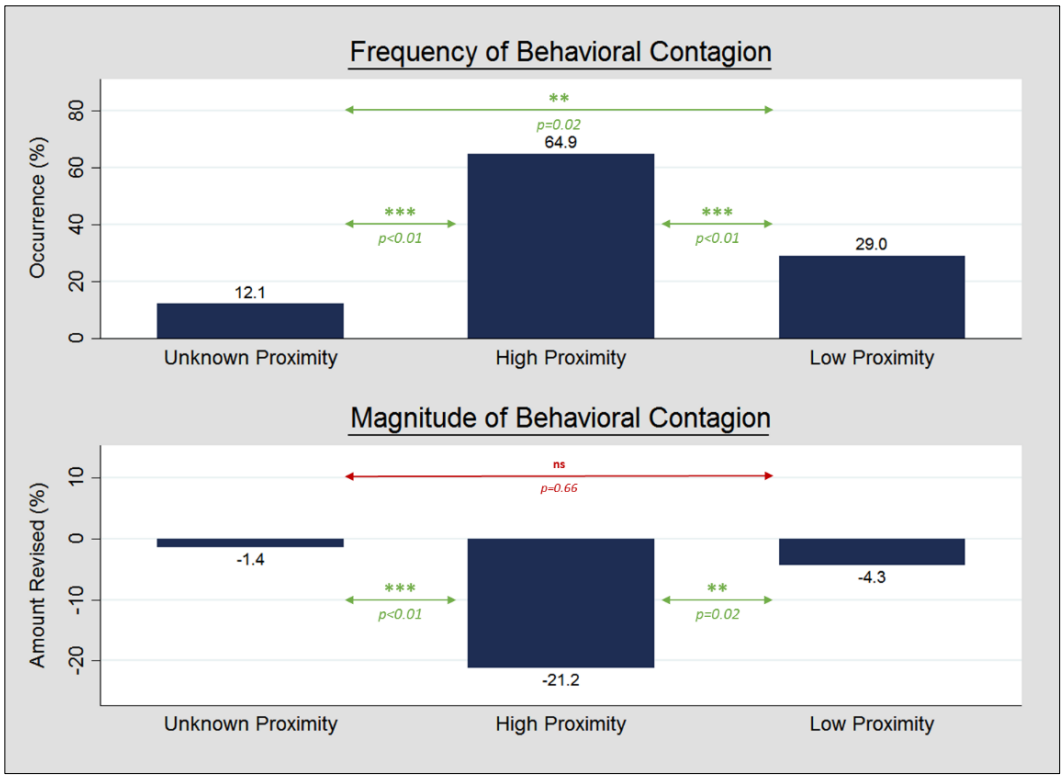

Figure 3: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on social proximity to peer. Definitions of *frequency of behavioral contagion* and *magnitude of behavioral contagion* are identical to those used in Figure 2. Unknown / high / low proximity refers to the respective proximity treatment conditions and indicates that the observer observed behavior of a random passive participant with unknown / high / low proximity score.

4.3 Composite Effect of Observed Behavior and Social Proximity

In a next step, 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 below in Figure 4.

By looking at the *frequency* of contagion first, the results suggest that behavioral contagion is asymmetric only where social proximity is high, with the difference being highly significant according to the equality of proportion statistic (88.2% vs. 45.0%, $p<0.01$). Contagion in the other proximity conditions is less frequent and also not significantly different between the types of observed behavior.

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

different at the 1% level (-51.3 vs. 4.4%, MWU, p<0.01) and 5% level (-10.1% vs. 1.7%, MWU, p=0.011), respectively.

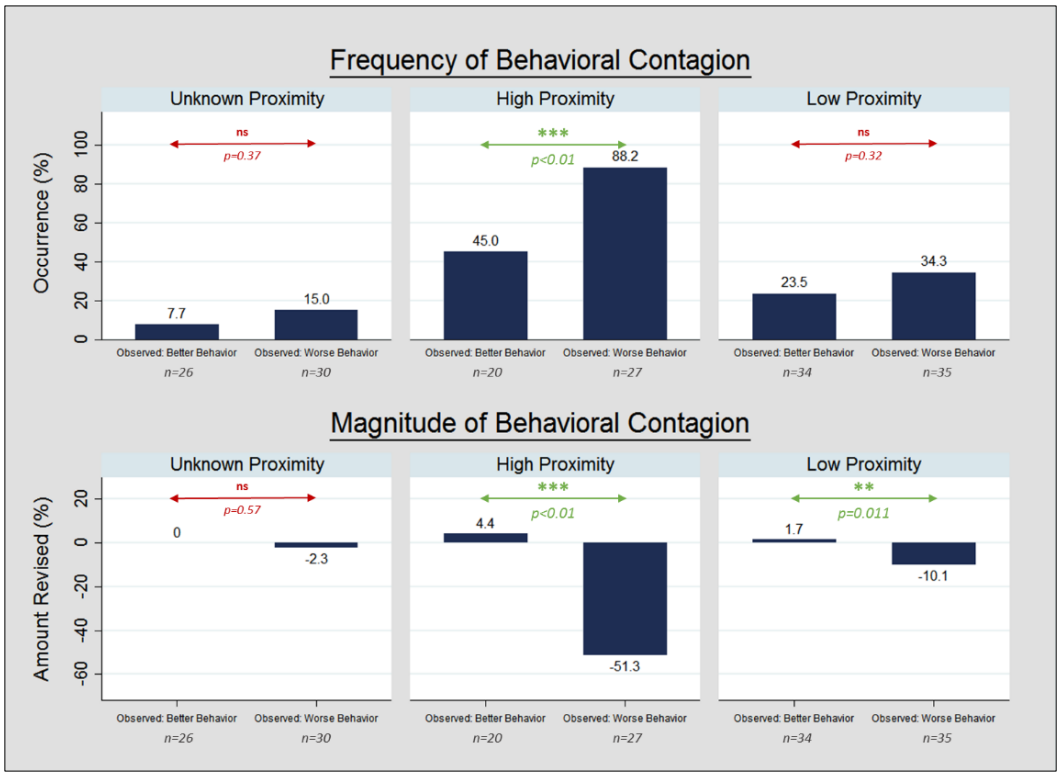

Figure 4: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on social proximity to peer. Definitions of frequency of behavioral contagion and magnitude of behavioral contagion are identical to those used in Figure 3. Definitions of observed behavior and proximity are identical to those used in Figures 3 and 3, respectively. A small number of participants observed exactly the same behavior as their own (and hence cannot be categorized as having observed better or worse behavior). For an analysis including these participants, see Appendix.

Again, our results are in support of our two hypotheses H_1 and H_2 and suggest that both anti-social behavior is more contagious than pro-social behavior and that social proximity amplifies the contagion of pro- and anti-social behavior, but does so more for anti-social behavior.[16]

[16] An alternative way of examining the data is to look at differences in behavioral contagion across different proximity specifications within the same type of observed behavior. The results indicate that both frequency and magnitude of behavioral contagion are significantly higher when behavior was observed in the high proximity condition, in particular when observing worse behavior, thus supporting the notion that anti-social behavior is overly contagious. We present the results in appendix Figure A2.

4.4 Whom Does Contagion Affect and How?

In a last step, we are interested in understanding whether and if so, to what extent behavioral contagion affects different individuals differently. In particular, we will analyze behavioral contagion conditional on the signaled type of the participants as in whether the participants have initially engaged in anti-social, pro-social or fair behavior. Conditional 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. 36.2%, p=0.02) or pro-social behavior (17.9% vs. 40.4%, p<0.01). The frequency of contagion in the latter two cases is statistically indistinguishable (36.2% vs. 40.4%, p=0.66).

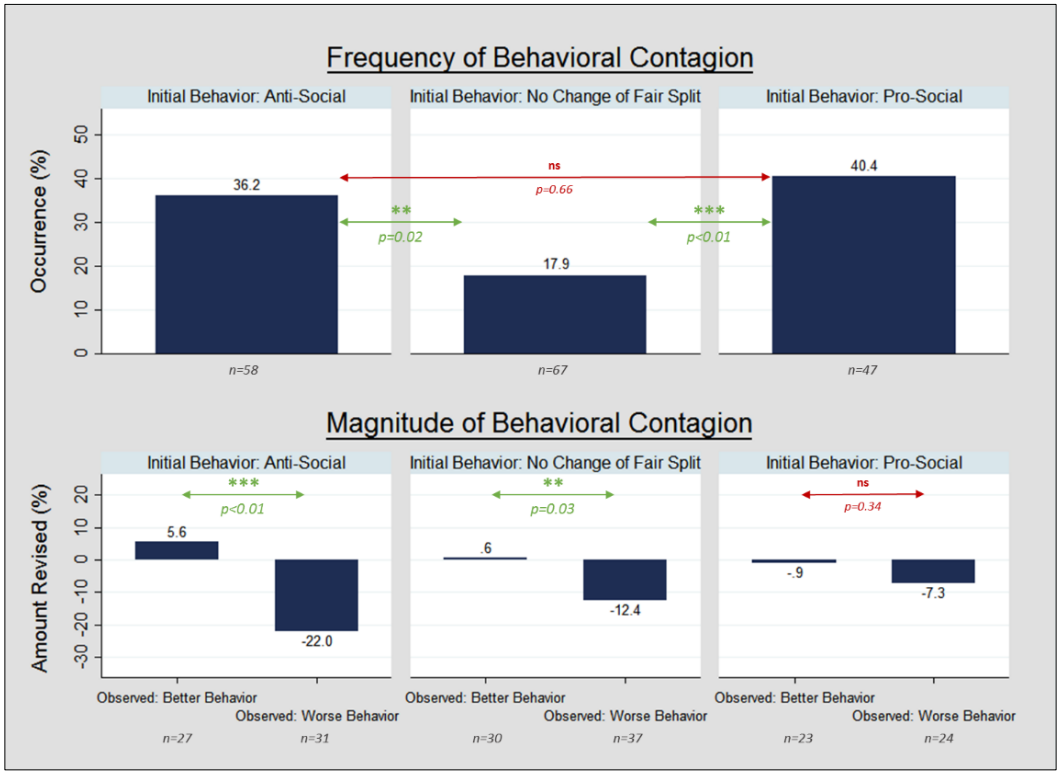

Figure 5: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on initial anti- and pro-social behavior. Definitions of *frequency of behavioral contagion* and *magnitude of behavioral contagion* are identical to those used in Figure 3. Definitions of observed behavior and proximity are identical to those used in Figures 2 and 3, respectively. A small number of participants observed exactly the same behavior as their own (and hence cannot be categorized as having observed better or worse behavior). For an analysis including these participants, see Appendix.

The pattern is even more 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 four times (= 22/5.6) as contagious as observing pro-social behavior, with the difference being statistically highly significant (-22.0% 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=0.02), but not for those participants who initially engaged in pro-social behavior (-7.3% vs. -0.9%, MWU, p=0.34).

In Figure 6, we slice the data even further and only look at initially anti-social (pro-social) participants who observed the opposite behavior. That is a pro-social (anti-social) peer, the results stay qualitatively the same and being 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.

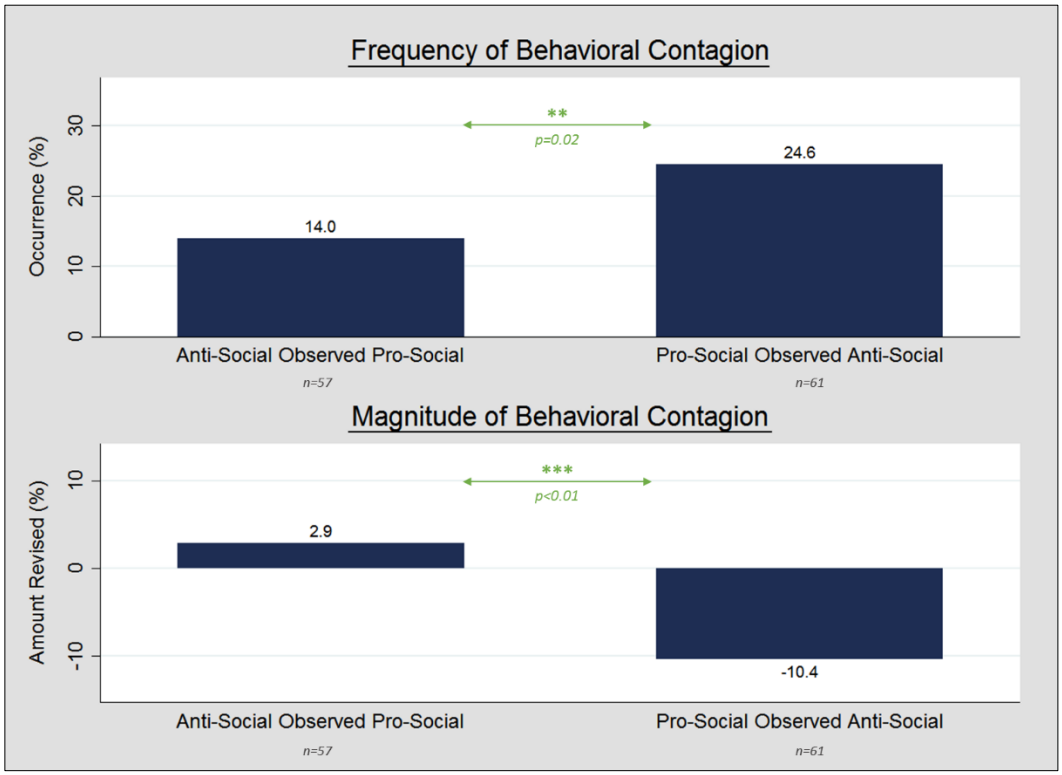

Figure 6: Frequency (upper section) and magnitude (lower section) of behavioral contagion conditional on observing the opposite behavior than one's own initial behavior. The graph depicts the behavioral contagion of initially pro-social (anti-social) participants who observed anti-social (pro-social) behavior. Definitions of *frequency of behavioral contagion* and *magnitude of behavioral contagion* are identical to those used in Figure 2. Definitions of observed behavior and proximity are identical to those used in Figures 2 and 3, respectively.

All these findings are in line with the idea that anti-social behavior carries an inherent ambiguity unless the norm is validated, i.e., via observation, while pro-social behavior is unconditional (for an elaborated discussion, see Bicchieri (2006)). These and the previous results support hypothesis $\mathbf{H_1}$ and confirm that anti-social behavior is more contagious than pro-social behavior.

We further investigate the robustness of our results using various regression specifications. We analyze the frequency (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, the observed anti- or pro-social behavior), an active participant's initial behavioral choice, the initial behavioral gap between the active and 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 a standardized measure of greed and risk, a number of interactions, as well dummies for the questions used in the personality questionnaire).

Our findings yield a number of interesting results and can be summarized as follows. We robustly find that higher peer proximity triggers more behavioral contagion, leading to substantially more anti-social behavior overall. What is more, the exposure to anti-social peer behavior begets more anti-social behavior, in particular for the already anti-social individuals in both frequency and magnitude. This group of people is also found to be particularly receptive to high proximity. In addition, we elicited incentivized beliefs about peer behavior prior to observation. As the regressions results suggest, such beliefs yielded no explanatory power, thus indicating that not the simple thought but action is driving the contagion of behavior. All remaining controls turn out not to be robustly significant across different specifications. These findings are in line with the non-parametric results presented before, indicating that behavioral contagion is particularly present in environments with high social proximity to one's peers, which overall leads to a crowding-out of pro-social behavior. We report the results in Table 2.

	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***	17.0398***	11.6685**	-25.2885***	-23.4287***	-31.1235**
	(21.6982)	(10.4019)	(12.3111)	(8.5234)	(7.8437)	(12.9994)
Low Proximity	4.2097**	3.3951**	4.1459	-4.4984	-4.7645	-6.2054
	(2.6540)	(1.7494)	(3.9636)	(4.7669)	(4.0026)	(6.5265)
Observed Anti-Social Behavior	4.4750***	1.5868	2.7879**	-18.4986***	-14.5667**	-17.7984***
	(2.4085)	(1.2008)	(1.2710)	(4.9563)	(6.9363)	(5.0044)
Initial Behavior						
(Base Level: No Change of Fair Split)						
Pro-Social	3.7867*	4.7251*	4.7238	8.7674	7.2860	2.8101
	(2.6295)	(4.1336)	(5.1273)	(6.8521)	(7.9709)	(10.9315)
Anti-Social	2.2859	0.3858	1.2898	1.2946	16.1079**	-4.3002
	(1.3900)	(0.4117)	(1.4976)	(4.1824)	(7.9567)	(5.6830)
Initial Behavioral Gap	0.9900*	1.0065	0.9864**	0.0143	-0.1123	0.0123
	(0.0058)	(0.0102)	(0.0059)	(0.0758)	(0.0847)	(0.0651)
Beliefs About Peer Behavior						
Pro-Social	0.5854	0.5609	0.6842	-7.2979	-5.9735	-5.5917
	(0.3672)	(0.3283)	(0.3934)	(6.0457)	(5.6881)	(5.5674)
Anti-Social	0.4791	0.4924	0.5473	-4.6446	-1.8528	-2.1482
	(0.2950)	(0.2663)	(0.2947)	(4.7348)	(4.3751)	(4.6539)
Gender	1.1204	1.0824	2.2918	-5.1159	-12.2865*	3.4026
	(0.7212)	(0.7653)	(2.2786)	(7.4739)	(6.3247)	(6.5263)
Greed		0.9602	0.9903		1.2555	1.2465
		(0.2089)	(0.2142)		(2.1982)	(2.1148)
Risk		0.8108	0.8692		1.6488	1.0497
		(0.1816)	(0.1962)		(2.2310)	(2.2754)
Interaction 1						
(Initial Behavior x Observed Behavior)						
Pro-Social & Anti-Social		0.5578			3.2578	
		(0.6998)			(12.2056)	
Anti-Social & Anti-Social		24.4830**			-26.0187**	
		(37.5465)			(13.0003)	
Interaction 2		0.7710			8.0022	
(Gender x Observed Behavior)		(0.6806)			(8.6649)	
Interaction 3						
(Initial Behavior x Social Proximity)						
Pro-Social & High Prox.			1.4476			25.6777
			(2.0599)			(19.3181)
Pro-Social & Low Prox.			1.1151			10.0039
			(1.2831)			(11.6835)
Anti-Social & High Prox.			5.9113			22.2268
			(8.6642)			(16.8720)
Anti-Social & Low Prox.			1.8578			4.4139
			(2.3124)			(8.4481)
Interaction 4						
(Initial Behavior x Gender)						
Pro-Social			0.4139			-9.1862
			(0.4325)			(11.8651)
Anti-Social			0.6707			-0.7228
			(0.7288)			(8.5743)
Interaction 5						
(Gender x Social Proximity)						
High Proximity			0.5251			-29.3372*
			(0.6534)			(16.6296)
Low Proximity			0.3944			-6.5315
			(0.4057)			(8.6067)
Dummies Questionnaire	Yes	No	No	Yes	No	No
Observations	172	172	172	172	172	172

Table 2: *Odds ratios (Logit) and regression coefficients (OLS). Standard errors in parentheses. *** Significant at the 1% level; ** Significant at the 5% level; * Significant at the 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 (= amount revised relative to one's initial behavior) for those participants who observed more anti- or pro-social behavior. 'Proximity' represents 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' depicts the active participant's initial decision and has the value 1 (2) if the participant initially decided to be anti- social (pro-social) by taking money away from (giving money to) charity. 'Initial behavioral gap' represents the difference in the active participant's initial behavior and the behavior he observed (in ECU) and takes on a theoretical value from -600 to +600 (given that the maximum anti- and pro-social behavior was -300 and +300, respectively). 'Beliefs about peer behavior' and 'Gender' are a dummy that have the 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. 'Interactions 1-5' represent the respective interactions and 'Dummies Questionnaire' indicate the inclusion of the items used in the initial personality questionnaire used to generate the proximity score.*

5. Discussion

Overall, we found convincing support for both of our hypotheses: anti-social behavior is indeed more contagious than pro-social behavior (H_1), and social proximity drives the magnitude of behavioral contagion, especially that of anti-social behavior (H_2). These 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.

While some of our results resonate with existing literature, some are generally new and provide an in-depth insight into the mechanisms of behavioral contagion of pro- and anti-social behavior. In real-life, the appropriateness of anti-social behavior is thought to be more ambiguous than that of pro-social behavior. The nature of anti-social or unethical behavior implies the overstepping of (social) boundaries or infringing laws. It can thus be expected that individuals often might want to engage in anti-social behavior but are first looking for social clues and signals to justify such behavior to themselves.

With reference to the different concepts in economics and (social) psychology explaining behavioral contagion and in light of our design, some theories are better at explaining behavior in our experiment than others are. Contagion observed in our experiment is likely to be explained by the theories of *social decisions and social distance* (Akerlof (1997), Glaeser & Scheinkman (2004)), *imitation of preferences* (Sliwka, 2007), as well as by some of the theories in (social) psychology, such as *social learning* (Bandura, 1971) and *norms* (Cialdini et al. (1990) and Bicchieri (2006)). Other theories are not applicable due to reasons of absence of learning better outcomes (which is required by Alós-Ferrer & Schlag's (2009) imitation concept), or the ability to not only observe but also be observed by the peers (which is required by, among others, Bernheim's (1994) theory on taste for conformity). We shall not attempt to disentangle which theory best explains behavioral contagion observed in our experiment but leave this venue for future research (see Dimant (2015) for a comprehensive discussion).

For the study of peer-effects in the lab this implies that, just as in the *real world*, individuals are often at the tipping point of behaving anti-socially but send a false signal (e.g.

behaving pro-socially or engage in neither behavior) until they come across a sufficiently salient justification device (e.g. norm-signaling through peer behavior). Research also indicates bad social norms to be particularly sticky, thus rendering it likely for behavioral contagion to be asymmetric towards the spillover of anti-social behavior, as observed in our experiment (Smerdon, Offerman & Gneezy (2016); Bicchieri, Dimant & Gächter (2017)). Such a carry-over of behavior from the *real world* to the lab serves a validation of our experimental approach (Kimbrough & Vostroknutov, 2016). In addition, the observed contagion is consistent with an individual's inherent desire for conformity as suggested by Bernheim (1994). The fact that such behavior exists even when individual behavior has no signaling value to the peers, as implemented in our experimental design by the unobservability of one's revised behavior, render Bernheim's findings even more general as previously suggested in literature, since behavioral contagion is found even when peer pressure to conform is absent.

What is more, 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. Irrespective of the existence of potential pre-existing norms, our design renders our main finding valid: behavioral contagion is highest where social coherence is strongest. If we follow the idea that the participants try to guess what the norm is, then the observation might get more weight when it is from a similar other. 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, in turn, would explain our results and the magnified impact of social proximity on contagion. Since we report lower-bound results, this is a strong indication that our method of inducing social proximity works and is likely to produce even stronger results when social proximity is introduced in a more sophisticated way.[17]

[17] We address potential experimenter demand concerns (for a discussion, see Zizzo (2010) in two ways. Firstly, consistent with the justification used by Shang & Croson (2008), whose methodology is closely related to ours, we argue that our setup enhances external validity because it mimics techniques that can be used in the field to raise money. Secondly, we reduce any potential experimental demand effect by being as transparent as possible in the experiment and do not withhold information from the participant with regards to either the matching algorithm or the treatment variation. That is, participants know exactly that there is a random but equal chance that they will end up in the unknown, low, or high proximity treatment and were given the particular differences between each treatment. It is brought to their attention that nature will draw this outcome and that the information is then presented accordingly. This transparency in combination with the fact that all treatments were always 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

6. Conclusion and Outlook

Deviant behavior that benefits oneself at the expense of others is both socially harmful and results in second-best solutions that are distortive from a welfare perspective. Conversely, the voluntary redistribution of money to those who have the least, e.g. in the form of donations, is socially desired. It is worthwhile to understand the underlying mechanism that drives pro- and anti-social behavior in order to implement effective policy measures.

Beyond pure self-maximizing considerations, behavior is also the result of social interactions in which conformity to particular behavior is affected by one's peers (Akerlof (1997), Glaeser & Scheinkman (2004)). One particular lens through which social interactions occur is through peer effects, which play a decisive role in explaining societal and economic outcomes. A battery of behavioral traits affects the shape and magnitude in which social interactions occur. An extensive stream of literature suggests that individuals are social animals and care for esteem, respect, and reputation (cf. Akerlof & Kranton (2000), Charness & Rabin (2002), Bénabou & Tirole (2011)). As such, individuals steadily act in social environments that define their role and standing within the social group. They do so by solving their inner rift between own desires and peer behavior through adjustment of identities and actions (Oxoby (2004), Akerlof (2017), Kranton (2016)).

Although peer effects have been extensively studied in different contexts both in the field and in the lab, research is still at an early understanding of the role of peer effects in pro- and anti-social settings. It is this paper's goal 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. Scholars across various fields are still at odds about whether the methods used qualify to observe clean peer effects, or whether a large part of our observations are an artifact of potential experimental confounds. A stream of literature points at methodological problems in soundly measuring such effects, especially outside the controlled laboratory environment (for a critical discussion see Manski (1993) & (2000) and Angrist (2014)). Our work addresses this criticism by the use of a novel design that allows us to

additional information willingly implemented by the experimenter to trigger any particular behavior. Comments made by Edward Glaeser, Daniel Houser, Lawrence 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.

study different types of behavior and behavioral spillovers. Being aware of the naturally occurring difficulty in studying social coherence in an artificial but controlled laboratory setting without introducing potential confounds, we capitalize on a novel approach to mimic social proximity as well as measure and vary social distance between peers.

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 spillovers, and answers the following two questions in particular: 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 coherence with one's peers in affecting behavioral contagion? For our purposes, we extend a variant of the dictator game as introduced by List (2007) and Bardsley (2008), among others, in which the dictator can give or take money. To provide clean evidence on peer effects, we capitalize on a one-shot dictator game in which participants are given the opportunity to give money to or take money away from the charity before and after learning peer behavior. Treatment variations include different levels of social proximity to the observed peers, which we incept by the novel design using dating website questions. This approach opens the venue for a broad range of future research examining the role of social proximity in coop-eration, reciprocity, and punishment behavior.

Our results suggest that anti-social behavior is more contagious than pro-social behavior, especially concerning initially anti-social and fair individuals, and the extent of social proximity with one's peers matters more for anti-social than for pro-social behavioral contagion. 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 salient signal that it is socially acceptable to be selfish. In light of the very conservative implementation of social proximity, we deem these results to represent a lower bound, thus strengthening the role of social proximity within the context of behavioral contagion.

In our experimental setting, we make use of a charity for the sole purpose of making pro- and anti-social behavior salient to the participant, which allows us to draw more general policy implication than only those restricted to charitable giving. Understanding social interactions in general and in particular the potentially resulting peer effects is fundamental from a

policy perspective. It does not only help to understand societal and economic outcomes beyond what standard economic forces can explain, but it also allows us to implement better-targeted policy measures to tackle a battery of challenges such as reducing crime rates, improving health conditions, or increasing labor market participation. Our results stress that social proximity is helpful in changing the behavior of anti-social people in particular. Exemplarily, in addition to providing information about other peer's behavior the respective companies could provide overlapping social proximity information in order to facilitate to facilitate donations to charities, increase towel reuse in hotels or the like. Our results suggest that this not only increases the saliency of observed behavior but also is more likely to trigger behavioral changes.

We deliberately refrained from applying sophisticated measures of any kind to induce social coherence. Rather, we resort to an easy-to-use methodological approach that may be used in future experiments in which inducing and measuring salient varying social proximity is key. Our lower-bound approach comes from the fact that participants were neither told the exact matching percentage nor the actual interests and preferences they had in common. In the treatments that included social proximity information, participants only knew whether they were observing a peer with above or below average congruence. A more sophisticated way to induce social proximity and match participants accordingly is likely to produce results that are more pronounced. The well-engineered mechanisms implemented by dating websites to match people and achieve high success rates are a shining example of what is possible: excluding matching partners based on personality traits that represent a no-go (e.g. smoking), putting emphasis on particular interests (e.g. sports), or individual characteristics (e.g. looks).

Our work may potentially add to a gradual move towards understanding the micro-foundations of social coherence. The experiment illustrates how individuals solve the inner rift between conformation and differentiation in a setting in which the saliency of social norms is mediated by social coherence, which helps to explain our result that proximity is the predominant driving behavior (Oxoby (2004), Akerlof (2017), Kranton (2016)). Much more scientific research is required to generate reliable policy measures to achieve both, more pro-social and less anti-social behavior. Exemplarily, recent MTO-research points at gender differences in behavioral assimilation (cf. Chetty, Hendren & Katz (2016)). Our research also helps to better

understand the "bad apple" model of peer effects and inform policy interventions in for example education and work environments, where competition and striving for status are a trigger for deviant behavior (cf. Charness, Masclet & Villeval (2013), Imberman, Kugler & Sacerdcote (2012)). There, isolated initial bad behavior that can spread uncontrolledly like wildfire, as can be observed in the recent Wells Fargo and Volkswagen scandals. It is important to understand to which extent these differences are driven by the two factors studied in this paper: the pro-/anti-sociality of peer behavior one is exposed to and the magnitude of social coherence with one's peers.

Shedding more light on the particular channels through which behavioral contagion emerges and their relative importance to one another (e.g. norms versus inherent preference for conformity) will help us better understand the underlying behavioral mechanisms and facilitate successful nudges. Our research is hopefully one of many more contributions to come on this matter.

7. Bibliography

Akerlof, G., 1997. Social Distance and Social Decisions. *Econometrica*, 65(5), pp. 1005-1027.

Akerlof, G. & Kranton, R., 2000. Economics and Identity. *Quarterly Journal of Economics*, 115(3), pp. 715-753.

Akerlof, R., 2017. Value formation: the role of esteem. *Games and Economic Behavior*, Volume 102, pp. 1-19.

Alos-Ferrer, C. & Schlag, K., 2009. Imitation and Learning. In: P. Anand, P. Pattanaik & C. Puppe, eds. *The Handbook of Rational and Social Choice.* Oxford University Press.

Andreoni, J., 1995. Warm-Glow Versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments. *Quarterly Journal of Economics*, 110(1), pp. 1-21.

Angrist, J. D., 2014. The perils of peer effects. *Labour Economics*, Volume 30, pp. 98-108.

Ariely, D., 2013. *The (Honest) Truth About Dishonesty: How We Lie to Everyone - Especially Ourselves.* Harper Collins Publ. UK.

Aronson, E., 2011. *The Social Animal.* 11 ed. Worth Publishers.

Azmat, G. & Iriberri, N., 2010. The Importance of Relative Performance Feedback Information: Evidence from a Natural Experiment Using High School Students. *Journal of Public Economics*, Volume 94, pp. 435-452.

Baccara, M. & Yariv, L., 2013. Homophily in Peer Groups. *American Economic Journal: Microeconomics*, 5(3), pp. 69-96.

Bandiera, O., Barankay, I. & Rasul, I., 2005. Social Preferences and the Response to Incentives: Evidence from Personnel Data. *The Quarterly Journal of Economics*, 120(3), pp. 917-962.

Bandiera, O. & Rasul, I., 2006. Social Networks and Technology Adoption in Northern Mozambique. *The Economic Journal*, 116(514), pp. 869-902.

Bandura, A., 1971. *Social Learning Theory.* New York: General Learning Press.

Bardsley, N., 2008. Dictator game giving: altruism or artefact. *Experimental Economics*, Volume 11, pp. 122-133.

Bénabou, R. & Tirole, J., 2011. Identity, morals, and taboos: beliefs as assets. *Quarterly Journal of Economics*, 126(2), pp. 805-855.

Benjamin, D., Choi, J. & Strickland, A., 2010. Social Identity and Preferences. *American Economic Review*, 100(4), pp. 1913-1928.

Bernheim, B. D., 1994. A Theory of Conformity. *Journal of Political Economy*, 102(5), pp. 841-877.

Bicchieri, C., 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms.* New York: Cambridge University Press.

Bicchieri, C., Dimant, E. & Gächter, S., 2017. Is the Dark Side More Powerful After All? Experimental Evidence on the Battle of Norms. *Mimeo*.

Bicchieri, C. & Xiao, E., 2009. Do the Right Thing: But Only if Others do so. *Journal of Behavioral Decision Making*, 22(2), pp. 191-208.

Bikhchandani, S., Hirshleifer, D. & Welch, I., 1998. Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *The Journal of Economic Perspectives*, 12(3), p. 151–170.

Bohnet, I. & Frey, B. S., 1999. Social Distance and other-regarding Behavior in Dictator Games: Comment. *The American Economic Review*, 89(1), pp. 335-339.

Bursztyn, L., Ederer, F., Ferman, B. & Yuchtman, N., 2014. Understanding Mechanisms Underlying Peer Effects: Evidence From a Field Experiment on Financial Decisions. *Econometrica*, 82(4), pp. 1273-1301.

Carrillo, J. & Mariotti, T., 2000. Strategic ignorance as a self-disciplining device. *Review of Economic Studies*, Volume 67, pp. 529-544.

Case, A. & Katz, L., 1991. *The Company You Keep: The Effects of Family and Neighborhood on Disadvantaged Youths*, NBER Working Paper No. 3705.

Centola, D. & Macy, M., 2007. Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3), pp. 702-734.

Charness, G. & Gneezy, U., 2008. What's in a name? Anonymity and social distance in dictator and ultimatum games. *Journal of Economic Behavior & Organization*, Volume 68, pp. 29-35.

Charness, G., Haruvy, E. & Sonsino, D., 2007. Social distance and reciprocity: An Internet experiment. *Journal of Economic Behavior and Organization*, Volume 63, pp. 88-103.

Charness, G., Masclet, D. & Villeval, M. C., 2013. The Dark Side of Competition for Status. *Management Science*, 60(1), pp. 38-55.

Charness, G. & Rabin, M., 2002. Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics*, Volume 117, pp. 817-869.

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

Chen, Y., Harper, F., Konstan, J. & Xin Li, S., 2010. Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review*, 100(4), pp. 1358-1398.

Chen, Y. & Li, S., 2009. Group Identity and Social Preferences. *American Economic Review*, 99(1), pp. 431-457.

Chetty, R., Hendren, N. & Katz, L. F., 2016. The Effects of Exposure to Better Neighborhoods on Children: New Evidence From the Moving To Opportunity Experiment. *American Economic Review*, 106(4), pp. 855-902.

Cialdini, R., Reno, R. & Kallgren, C., 1990. A Focus Theory Of Normative Conduct: Recycling The Concept Of Norms To Reduce Littering in Public Places. *Journal of Personality and Social Psychology*, pp. 1015-1026.

Costa-Font, J. & Cowell, F., 2015. Social Identity and Redistributive Preferences: A Survey. *Journal of Economic Surveys*, 29(2), pp. 357-374.

Croson, R. & Shang, J., 2008. The Impact of Downward Social Information on Contribution Decisions. *Experimental Economics*, Volume 11, pp. 221-233.

Damm, A. P. & Dustmann, C., 2014. Does Growing Up in a High Crime Neighborhood Affect Youth Criminal Behavior?. *American Economic Review*, 104(6), pp. 1806-1832.

Dana, J., Weber, R. A. & Kuang, J. X., 2007. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, Volume 33, pp. 67-80.

DellaVigna, S., List, J. & Malmendier, U., 2012. Testing for altruism and social pressure in charitable giving. *Quarterly Journal of Economics*, 127(1), pp. 1-56.

Dimant, E., 2015. *On Peer Effects: Behavioral Contagion of (Un)Ethical Behavior and the Role of Social Identity*, MPRA Paper No. 68732.

Duflo, E., Dupas, P. & Kremer, M., 2011. Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya. *American Economic Review*, 101(5), pp. 1739-1774.

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

Eckel, C. & Petrie, R., 2011. Face Value. *American Economic Review*, 101(4), pp. 1497-1513.

Ellingsen, T. & Johannesson, M., 2008. Pride and prejudice: the human side of incentive theory. *American Economic Review*, 98(3), pp. 990-1008.

Falk, A. & Fischbacher, U., 2002. "Crime" in the Lab - Detecting Social Interaction. *European Economic Review*, Volume 46, pp. 859-869.

Falk, A. & Ichino, A., 2006. Clean Evidence on Peer Pressure. *Journal of Labor Economics*, Volume 24, pp. 39-57.

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

Frey, B. & Meier, S., 2004. Social Comparison and Pro-Social Behavior: Testing Conditional Cooperation in a Field Experiment. *American Economic Review*, 94(5), p. 1717–1722.

Gächter, S., Molleman, L. & Nosenzo, D., 2017. *Peer Effects in Norm Compliance,* Mimeo.

Gächter, S., Nosenzo, D. & Sefton, M., 2013. Peer Effects in Pro-Social Behavior: Social Norms or Social Preferences. *Journal of the European Economic Association*, 11(3), pp. 548-573.

Gino, F., Ayal, S. & Ariely, D., 2009. Contagion and Differentiation in Unethical Behavior. *Psychological Science*, 20(3), pp. 393-398.

Gino, F. & Bazerman, M. H., 2009. When misconduct goes unnoticed: the acceptability of gradual erosion in others' unethical behavior. *Journal of Experimental Social Psychology*, Volume 45, p. 708–719.

Gino, F. & Moore, D. A., 2007. Effect of Task Difficulty on Use of Advice. *Journal of Behavioral Decision Making*, 20(1), pp. 21-35.

Glaeser, E. L. & Scheinkman, J. A., 2004. Measuring Social Interactions. In: S. N. Durlauf & H. P. Young, eds. *Social Dynamics (Economic Learning and Social Evolution Series).* MIT University Press Group Ltd., pp. 83-131.

Gneezy, U., 2005. Deception: the role of consequences. *American Economic Review*, Volume 95, pp. 384-394.

Gould, E. & Kaplan, T., 2011. Learning unethical practices from a co-worker: the peer effect of Jose Canseco. *Labour Economics*, 18(3), pp. 338-348.

Greiner, B., 2004. An online recruitment system for economic experiments. *Forschung und Wissenschaftliches Rechnen*, Volume 63, pp. 79-93.

Grossman, Z., 2014. Strategic ignorance and the robustness of social preferences. *Management Science*, 60(11), pp. 2659-2665.

Herbst, D. & Mas, A., 2015. Peer effects on worker output in the laboratory generalize to the field. *Science*, 350(6260), pp. 545-549.

Hitsch, G., Hortaçsu, A. & Ariely, D., 2010. Matching and sorting in online dating. *American Economic Review*, 100(1), pp. 130-163.

Ichino, A. & Maggi, G., 2000. Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm. *The Quarterly Journal of Economics*, 115(3), pp. 1057-1090.

Imberman, S. A., Kugler, A. D. & Sacerdote, B. I., 2012. Katrina's Children: Evidence on the Structure of Peer Effects from Hurricane Evacuees. *American Economic Review*, 102(5), pp. 2048-2082.

Innes, R. & Mitra, A., 2013. Is dishonesty contagious?. *Economic Inquiry*, 51(1), pp. 722-734.

Kandel, E. & Lazear, E., 1992. Peer pressure and partnerships. *Journal of Political Economy*, Volume 100, pp. 801-817.

Katz, L. F., Kling, J. R. & Liebman, J. B., 2001. Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment. *The Quarterly Journal of Economics*, 116(2), pp. 607-654.

Kessler, J. & Milkman, K., forthcoming. Identity in Charitable Giving. *Management Science*.

Kimbrough, E. & Vostroknutov, A., 2016. Norms make preferences social. *Journal of the European Economic Association*, Volume 14, pp. 608-638.

Kling, J. R., Ludwig, J. & 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), pp. 87-130.

Kranton, R. E., 2016. Identity Economics 2016: Where do social distinctions and norms come from?. *American Economic Review: Papers & Proceedings*, 106(5), pp. 405-409.

Lavy, V. & Schlosser, A., 2011. Mechanisms and Impacts of Gender Peer Effects at School. *American Economic Journal: Applied Economics*, Volume 3, pp. 1-33.

Leider, S., Möbius, M. M., Rosenblat, T. & Do, Q.-A., 2009. Directed altruism and enforced reciprocity in social networks. *The Quarterly Journal of Economics*, 124(4), pp. 1815-1851.

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

Manski, C., 1993. Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies*, Volume 60, pp. 531-542.

Manski, C., 2000. Economic Analysis of Social Interactions. *The Journal of Economic Perspectives*, 14(3), pp. 115-136.

Mas, A. & Moretti, E., 2009. Peers at Work. *American Economic Review*, 99(1), pp. 112-145.

Offerman, T., 2002. Hurting hurts more than helping helps. *European Economic Review*, Volume 46, pp. 1423-1437.

Oxoby, R., 2004. Cognitive dissonance, status and growth of the underclass. *Economic Journal*, 114(498), pp. 727-749.

Sacerdote, B., 2001. Peer Effects with Random Assignment: Results for Dartmouth Roommates. *The Quarterly Journal of Economics Economics*, 116(2), pp. 681-704.

Shang, J. & 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), pp. 1422-1439.

Sliwka, D., 2007. Trust as a Signal of a Social Norm and the Hidden Costs of Incentive Schemes. *American Economic Review*, 97(3), pp. 999-1012.

Smerdon, D., Offerman, T. & Gneezy, U., 2016. *"Everybody's doing it": On the emergence and persistence of bad social norms*, Tinbergen Institute Discussion Paper: TI 2016-023/I.

Thöni, C. & Gächter, S., 2015. Peer Effects and Social Preferences in Voluntary Cooperation. *Journal of Economic Psychology*, Volume 48, pp. 72-88.

Wheeler, L., 1966. Toward a Theory of Behavioral Contagion. *Psychological Review*, 73(2), pp. 179-192.

Zizzo, D. J., 2010. Experimenter demand effects in economic experiments. *Experimental Economics*, Volume 13, pp. 75-98.

8. Appendix

Section A: Additional Figures and Robustness Checks

Figure A2 details our behavioral data in two ways: before observation in pooled form and after observation by treatment. Starting with a descriptive observation, we observed some heterogeneity across treatments in terms of the distribution of pro-/anti-social behavior. The horizontal axis depicts amount of money allocated to the charity. Here, 0% represents the equal *a priori* distribution of money between the participant (300 ECU) and the charity (300 ECU), i.e. a positive (negative) value indicates that the participant gave to (took from) the charity.

Our data indicates that prior to observing their peer's behavior, active participants' decisions mainly clustered around 300 ECU (0%), which represents the decision to not change the initial equal distribution between oneself and the charity. After the observation, however, we observe a perceptible skewness 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.

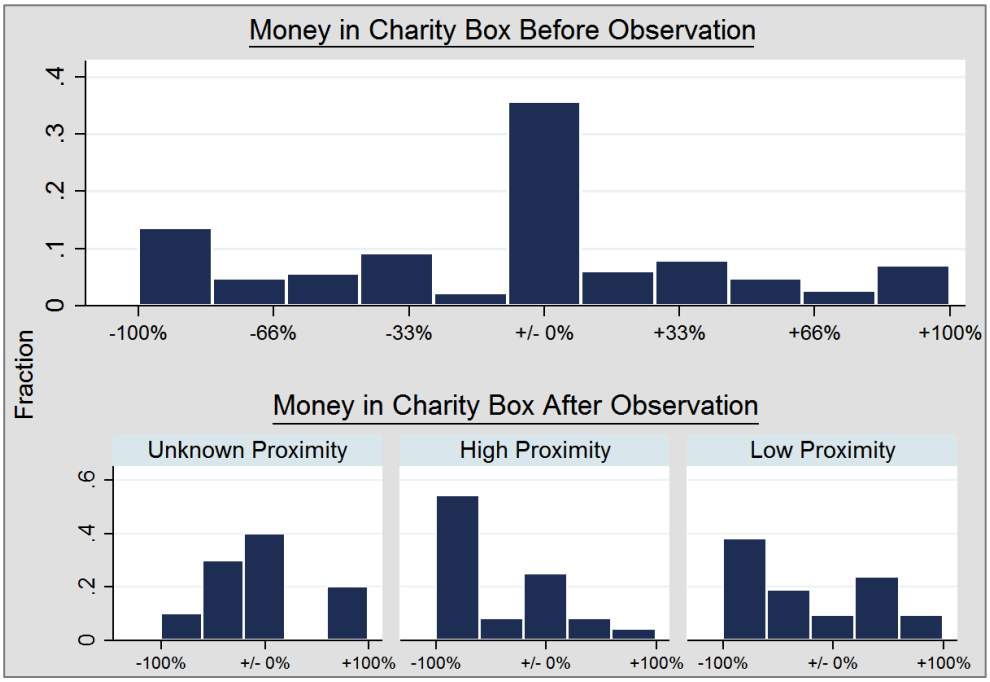

Figure A1: *Incidence of choices to not change the initial equal distribution / give money to / take money from the charity before and after peer observation. The horizontal axis depicts a continuum of money left in the respective charity's cash account, with 0% representing the status-quo allocation.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

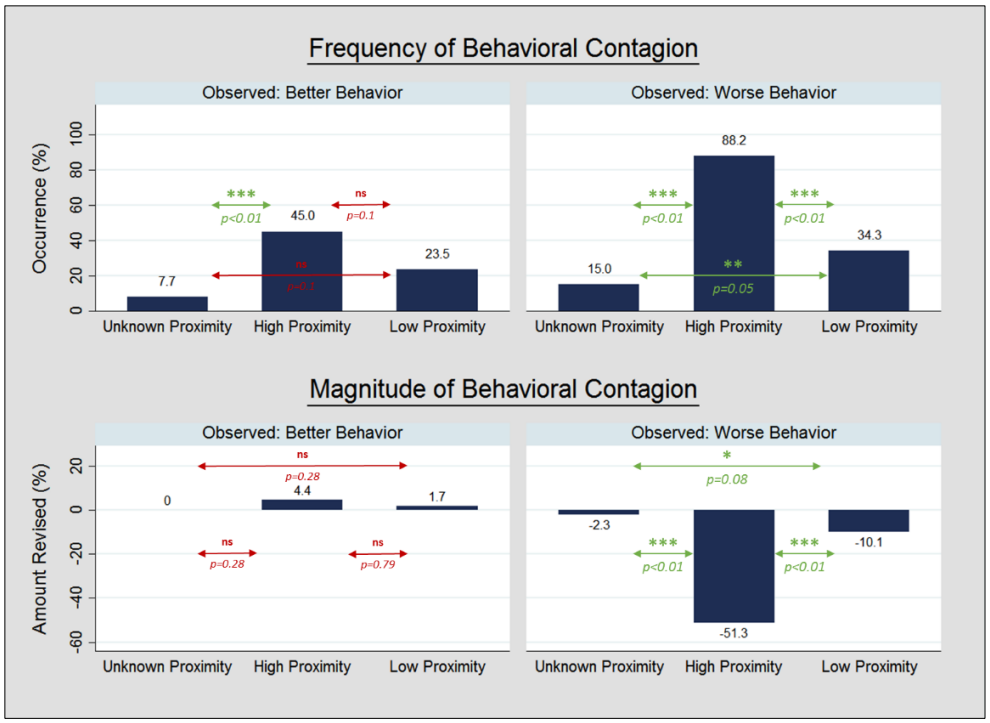

Figure A2: Alternative analysis to the composite effect of observed behavior and social proximity. See footnote 18 on page 19.

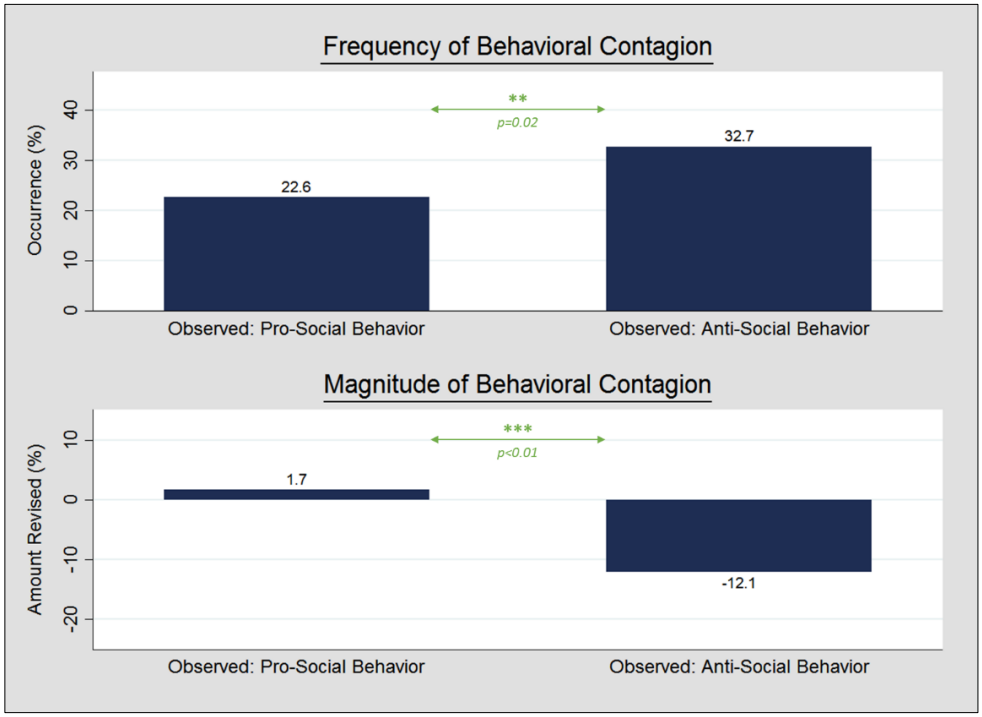

Figure A3: Robustness check of Figure 3. Frequency and magnitude of behavioral contagion by observation of pro- and anti-social behavior.

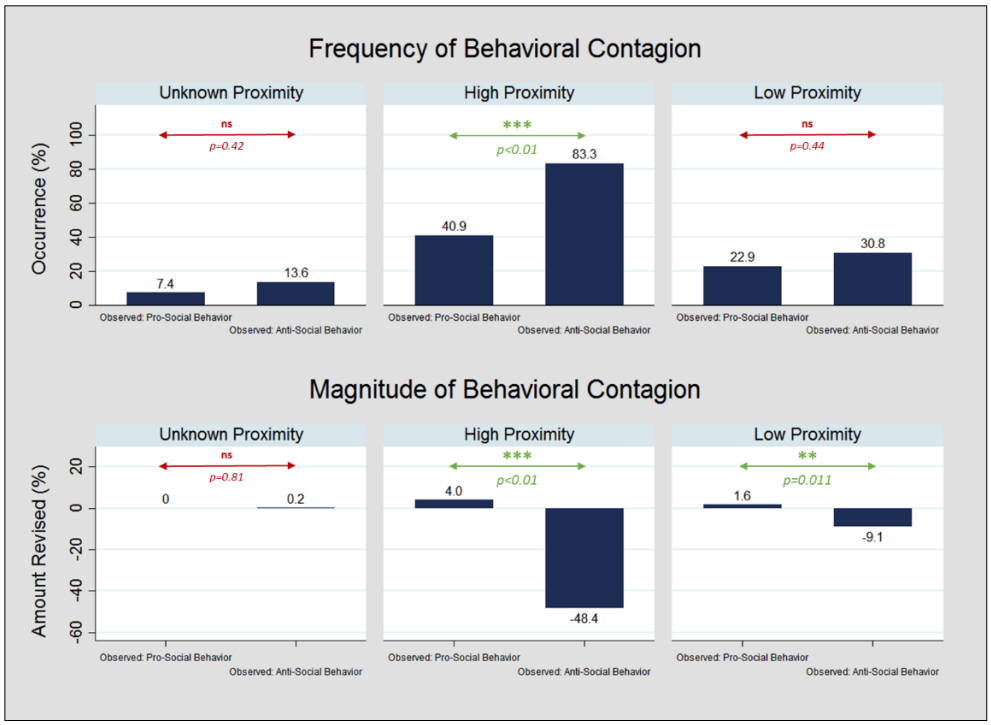

Figure A4: Robustness check of Figure 5. Frequency and magnitude of behavioral contagion by social proximity.

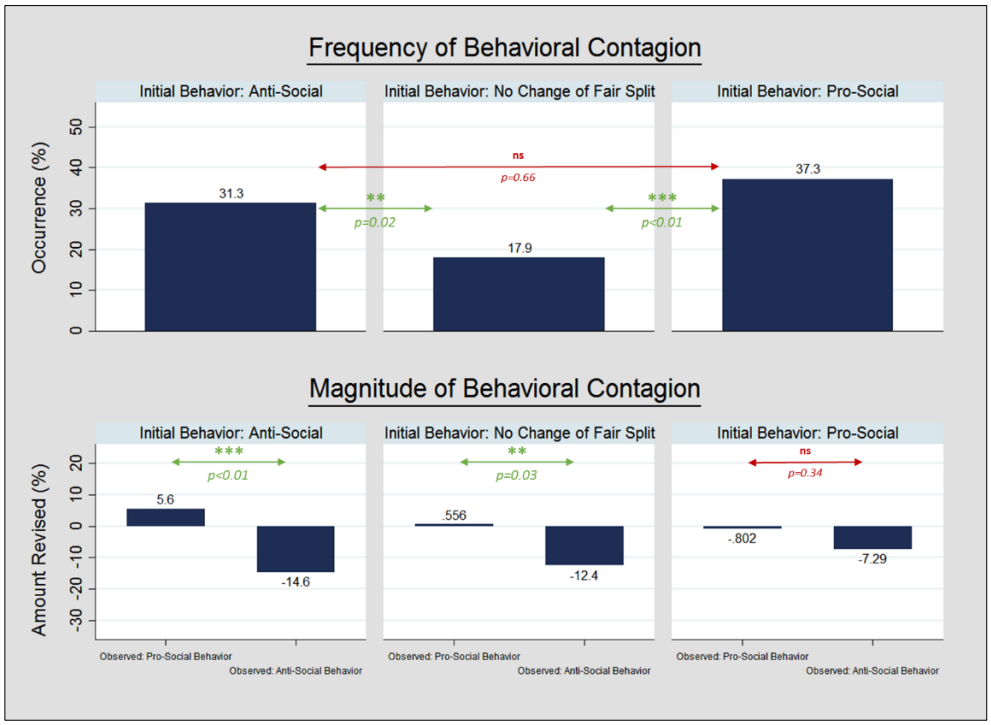

Figure A5: Robustness check of Figure 6. Frequency and magnitude of behavioral contagion by initial behavior, observed anti-/pro-social behavior and social proximity.

Section B: 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 "Taler" you earn during the experiment.
- At the end of the experiment, the amount of "Taler" that you have earned will be converted into real money at an exchange rate of 20 Taler = 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** as well as a **charitable organization of your choice** (i.e. an officially registered charity organization) will be provisionally assigned a monetary amount of **300 Taler** each.
- During the experiment you will have to decide on whether you want to…
  - … *take* a part or all of the money from the charitable organization.
  - … *leave* the division of the sum of money as it is.
  - … *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). In particular, we will double the respective amount and donate it to the charity after the experiment ends. 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 Taler** (including the show-up fee) at the end of the experiment.

- <u>The total payoff of the participants:</u>
  - **In case you are the randomly chosen participant**
    300 Taler +/- the amount of money that has been given to/taken from the cash account of the charitable organization
  - **In case you are <u>not</u> the randomly chosen participant**
    150 Taler

- <u>The total payoff of the charitable organization:</u>
  - (Amount of money in the cash account of the charitable organization of the randomly chosen participant) × 2
- 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

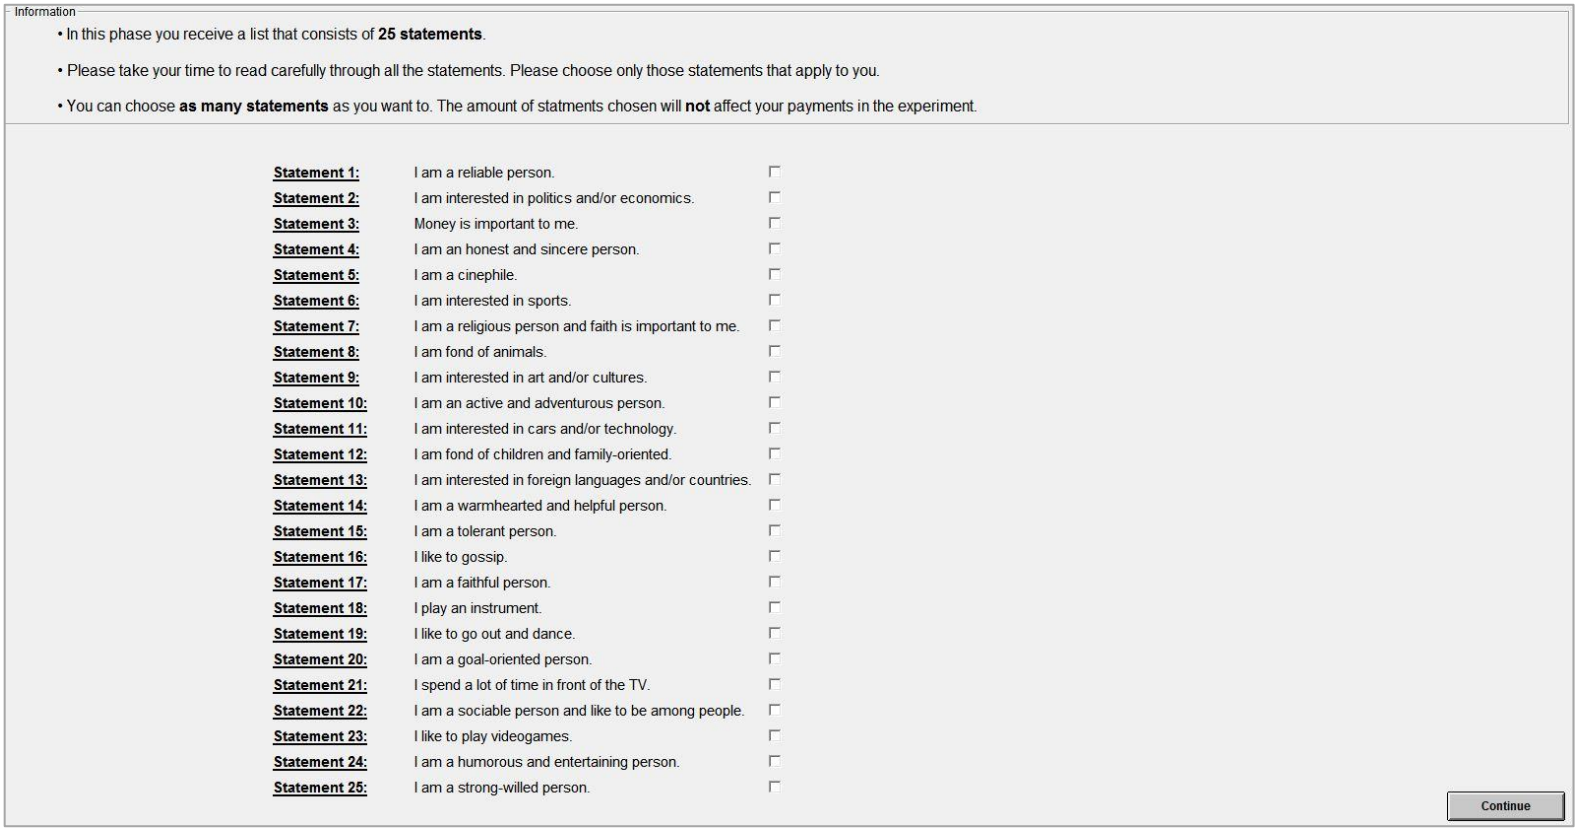

2. First decision towards charity

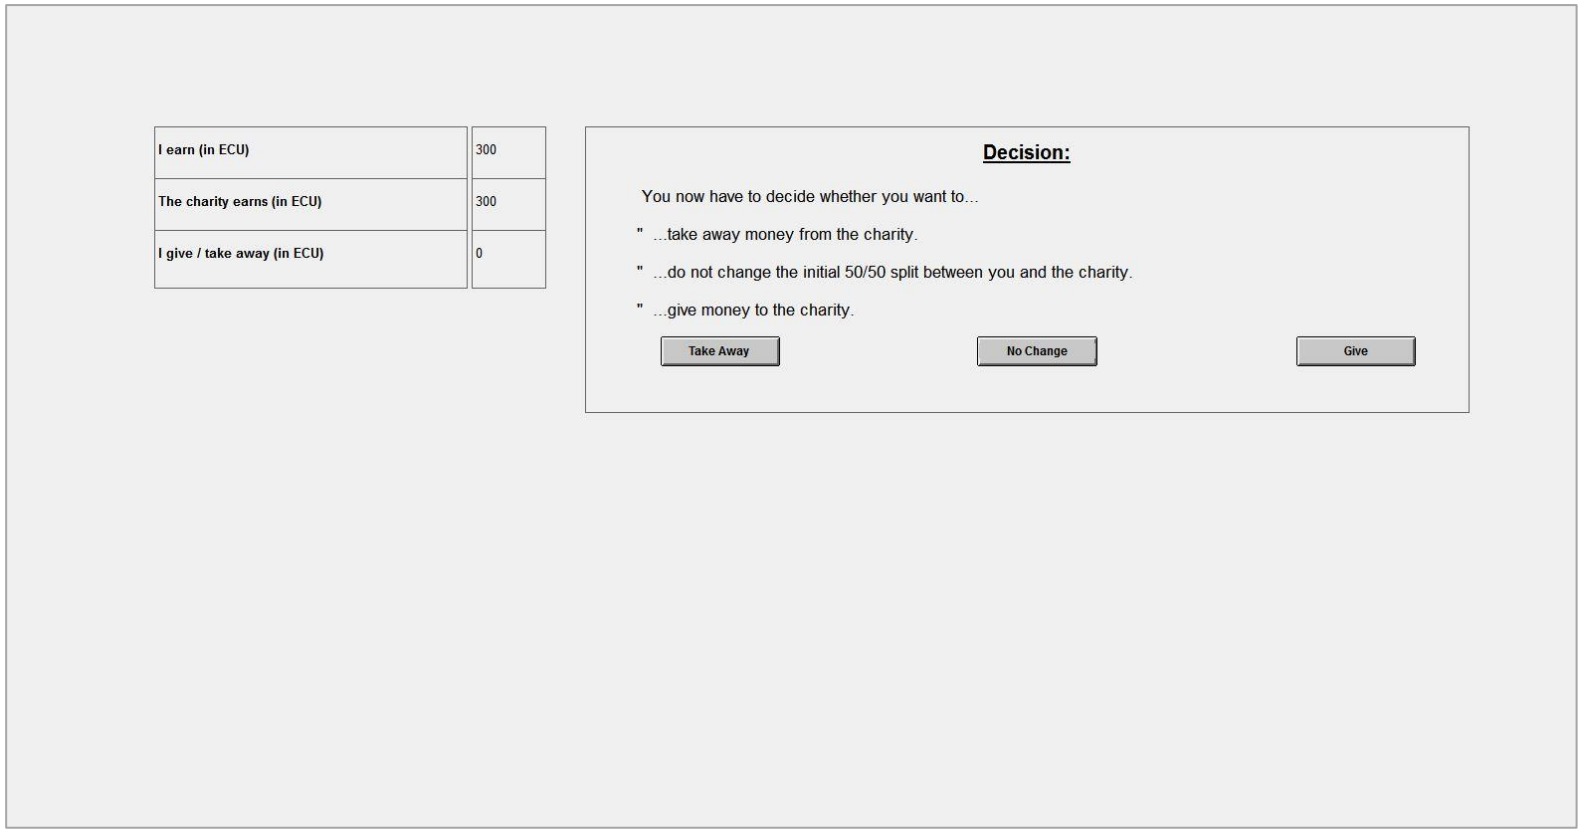

(Exemplarily for the taking away decision)

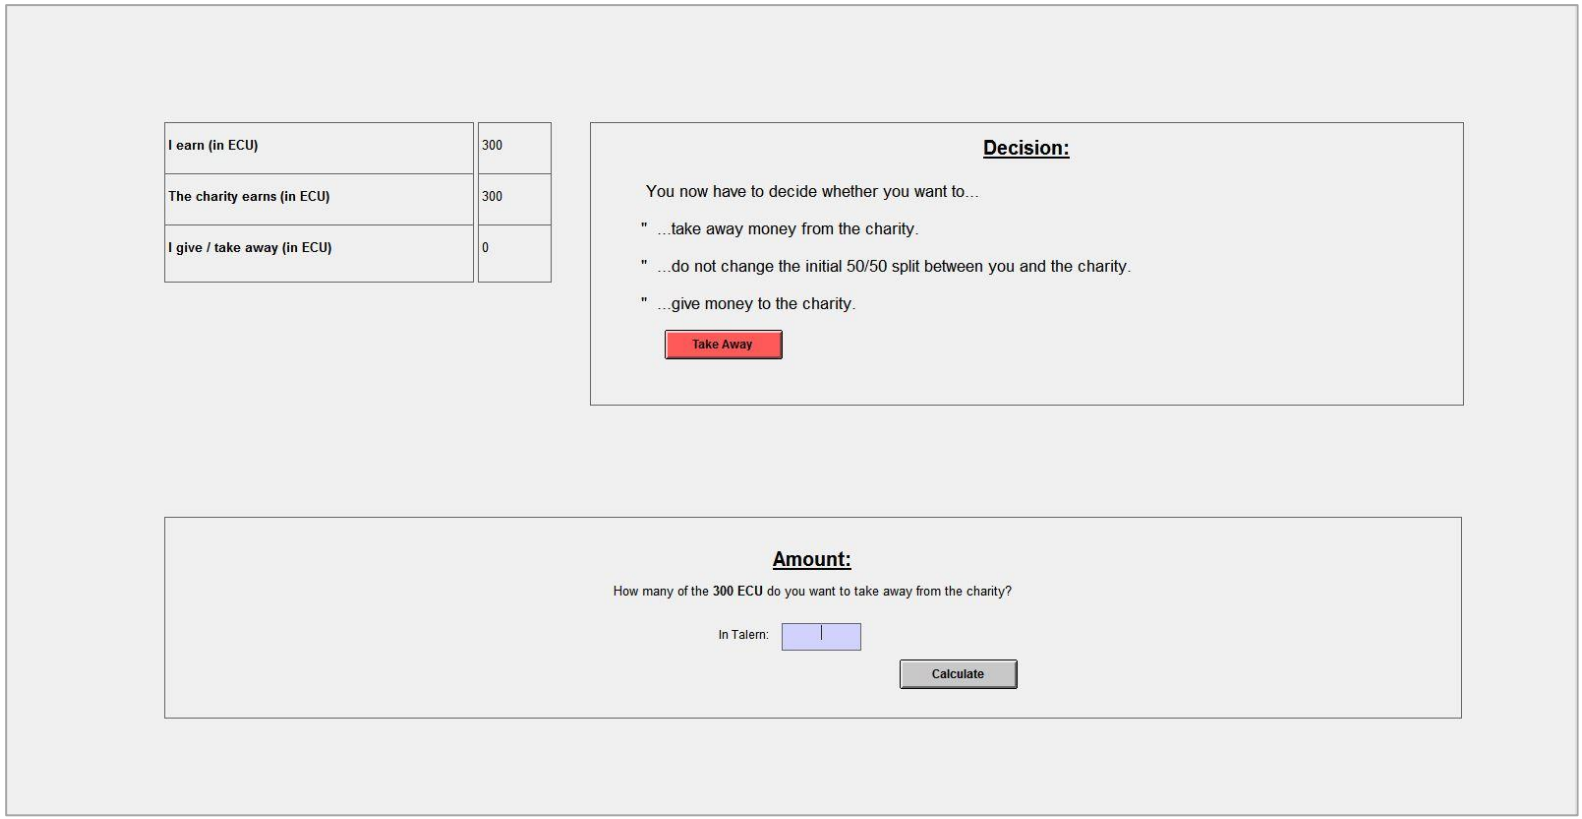

3. Observation of one's peer and potential revision of one's initial decision

Important: treatments vary by the information on the social proximity measure (first box): unknown similarity (Baseline), more similar (Treatment 1), less similar (Treatment 2).

(Exemplarily for revision of one's initial decision)