

Coordination and Contagion: Individual Connections and Peer Mechanisms in a Randomized Field Experiment*

Philip Babcock
University of California, Santa Barbara and NBER
babcock@econ.ucsb.edu

Kelly Bedard
University of California, Santa Barbara and IZA
kelly@econ.ucsb.edu

Stefanie Fischer
Cal Poly, San Luis Obispo and IZA
sjfische@calpoly.edu

John Hartman
University of California, Santa Barbara
hartman@econ.ucsb.edu

June 2019

Abstract

This paper investigates peer effects at the level of individual connections, leveraging the approach to shed light on peer mechanisms. In a field experiment using college freshmen, we elicited best friends and offered monetary incentives for gym visits to a treated subset. We find large spillovers from treated subjects to treated best friends but none from treated subjects to control best friends. We also find evidence of a mechanism: Subjects coordinate by visiting the gym with best friends, indicating that the intervention harnesses complementarities in utility or commitment mechanisms. Results highlight subtle peer effects and mechanisms that often go undetected.

* We thank Eric Bettinger, Olivier Deschenes, Peter Kuhn, Shelly Lundberg, Heather Royer, Bruce Sacerdote, Jon Sonstelie, Dick Startz, and Cathy Weinberger for helpful comments. We are grateful for funding by the Hellman Family Foundation, and excellent research assistance by Jennifer Carnan, Jennifer Schulte, Allison Nuovo, Christy Helvestine, Bonnie Queen, Jessica Evans, Natalie Brechtel, Randi Golde, Jennifer Milosch, Allison Bauer, and Ernesto Boffy-Rameriz. We are also grateful for assistance from Chris Clontz at the Recreation Center at the University of California, Santa Barbara. This experiment and research has been conducted with IRB approval from the University of California, Santa Barbara.

1 Introduction

The goal of this paper is to investigate peer effects at the level of individual connections—arguably, the level at which these effects are most likely to operate—and to leverage this approach to shed light on peer mechanisms. Previous work on peer effects has often relied on variation in peer reference groups that have been externally imposed or defined, rather than on self-identified friends. Since most real-world connections are endogenously formed, it is worth augmenting this research with studies of spillovers between individuals whose connections were self-chosen. This paper adds to previous research both because it penetrates the black box to identify peer mechanisms, and because it features exogenous variation in peer behavior while allowing peer effects to operate through chosen social connections.

We incentivize a random subset of college students by offering to pay them for repeated visits to the campus recreation center, as in Charness and Gneezy (2009), and use this carefully controlled setting to study spillovers and mechanisms. Much important research on peer effects has relied on random variation in assigned classes, grade-levels, roommates, residence halls, squadrons, or other imposed groupings,¹ rather than reference groups of self-chosen peers. Some studies of spillovers in fitness and obesity have focused on naturally occurring social connections, but these have tended to lack plausibly exogenous sources of variation. Christakis and Fowler (2007 and 2008), for example, argue that social spillovers in obesity are large, but their results have been criticized by Cohen-Cole and Fletcher (2008a and 2008b) due to the endogeneity of

¹ Examples include Sacerdote (2001), Zimmerman (2003), Boisjoly et al. (2006), Foster (2006), Lyle (2007), Kremer and Levy (2008), Carrell, Fullerton, and West (2009), Carrell, Hoekstra, and West (2011), Yakusheva et al. (2014), Fischer (2017) and Hoekstra et al. (2018).

friendship formations used in the study. An important feature of our experiment is that prior to treatment, we elicit subject's friends. Random assignment of treatment creates random variation in the treatment status of the peers to whom the subject is exposed, and from this we identify peer effects. The methodology we propose can be applied in other interventions, particularly for settings in which subjects have the potential to be connected to one another by naturally occurring social ties. These could include school-level interventions to raise test scores, promote college enrollment decisions, combat drug and alcohol abuse, or raise civic participation; workplace interventions to promote health or productivity; village-level interventions to improve sanitation, prevent disease, or raise school attendance; and other effort-elicitation contexts in which there is an opportunity for repeated social interaction over time. In many of these settings, it might be possible to use administrative data to identify important social connections before treatment assignment.

Our main findings are that treated subjects with treated best friends put forth significantly more effort toward the incentivized task than do treated subjects with control best friends. The peer effect is about 25 percent as large as the direct individual effect of the incentive. There is no observed spillover for control subjects, and thus a partial population experiment that lacked data on individual connections would not have isolated this peer effect. Further, we find evidence of a mechanism that explains the entire peer effect: subjects coordinate by visiting the gym jointly with their best friends, an indication that the intervention harnesses complementarities in utility or helps with commitment problems.

2 Background

Our strategy for investigating the well-known problems involved in identifying social effects is based largely on Moffitt's (2001) description of a partial population intervention. Partial population designs have recently been used to estimate spillovers hypothesized to arise from information transfer in schools or departments (Duflo and Saez, 2003; Miguel and Kremer, 2004), from imitation or social insurance between households (Angelucci and DiGiorgio, 2009; Kuhn et al., 2011; Bhattacharya 2018) and from learning or imitation in schooling decisions (Bobonis and Finan, 2009; Lalive and Cattaneo, 2009). One could conceptualize an ideal partial population design as one which begins with a sample of social groups or villages between which there are no spillovers. Within the groups or villages designated for treatment, a random subset of agents or households is treated. Differences in group-level means in this setting identify several quantities of interest: (a) The difference between mean outcomes for treated and control villages captures average group-level treatment effects; (b) The difference between mean outcomes of treated subjects in treated villages and control subjects in untreated villages estimates the average effect of village treatment on treated subjects; (c) The difference in the mean outcomes of untreated subjects in treatment and control villages captures the effect of village treatment on untreated subjects.

Though the partial population experiment is a powerful tool for introducing exogenous variation into settings with social spillovers, in practice there are several limitations associated with the design: (1) Agents are assumed to respond to mean behavior of the reference group, because the internal structure of the group is unknown. However, a subject is unlikely influenced by a peer he barely knows, even if both agents

are a part of the same reference group (e.g., a class, a grade, or a village). As such, when data on individual ties between individual agents is lacking, the reference group often includes irrelevant peers leading to reduced precision in the estimated peer effect. (2) The driving mechanism is often unclear. A spillover may be observed from treated subjects to controls, but there is often little that can be done to infer the mechanism. (3) Usually, only spillovers from treated subjects to controls are isolated and estimated.

To flesh out (3), we note that there are a number of concrete examples of partial population experiments that do not identify certain kinds of peer effects. In Kuhn et al. (2011), the authors estimate consumption spillovers on lottery non-participants resulting from exposure to lottery winners. But they cannot detect whether lottery winners in a winning postal code influence each other's consumption. Angelucci and DeGiorgi (2009) estimate consumption spillovers from households eligible for transfers to wealthier, ineligible households, but cannot estimate spillovers between eligible households. Lalive and Cattaneo (2009), using similar data, summarize concerns about precisely this kind of spillover between the treated (p. 460): "If children from poor households only interact with other children from poor households, there could be important social spillover effects that cannot be detected with the PROGRESA experiment." It bears emphasizing that the problem the authors describe is not unique to the PROGRESA experiment but is a general feature of partial population interventions that do not collect data on connections between individuals.

By incorporating social network data into the framework of a partial population experiment, we attempt to address these challenges. One advantage of our design is that we randomize at the individual level. This may be useful for several reasons. First, this is

likely where the most relevant action can be observed since individuals are influenced by a non-random subset of a group. It seems likely that individuals respond strongly to certain individuals in the group and are much less influenced by others. Second, data on friendship networks combined with treatment assignment at the individual level allow us to shed light on peer mechanisms. Third, our experiment will highlight spillovers that are present even though there is no discernible effect of group treatment on controls, i.e., *when the difference in group means described in (c) of an idealized partial population experiment is zero.*²

There are several previous studies that use friendship networks to examine spillovers in real world settings.³ Miguel and Kremer (2007) elicit friendship networks from some of the subjects in the intestinal worms intervention (those who were exposed to the program later). They examine how these ties predict subjects' decisions to take up treatment. Our paper resembles this work in the use of networks but differs in that our design is based on randomization at the individual level and that we elicit ties from all participating subjects. Similarly, Oster and Thornton (2012) elicit friendship ties and estimate spillovers in menstrual cup usage. There, the setting allows for the analysis to be simplified in a way that is not always practical in other environments: spillovers for the controls are not estimated (and may have been implausible, because controls were not provided with the product). Banerjee et al. (2013) collect household level network data in

² We note that it is possible, even without network data, to calculate spillovers from treated subjects to untreated subjects in a partial population experiment by running experiments in many different groups or villages and randomizing fraction treated (e.g., Philipson, 2000). However, network data allow the estimation of social spillovers using much less data and at less expense. If running experiments in enough villages to get statistical power is logistically challenging and expensive, or if the immediate social circle is likely to be a more salient reference group than the village, then an approach that relies on exogenous variation in treatment across individual social ties offers advantages.

³ In addition, there exists important research using real world social networks to explore altruism and reciprocity in experimental games. See Leider et al. (2009).

India prior to the entry of a microfinance institution to study the diffusion of microfinance loans where their model can distinguish between two mechanisms: information passing and endorsement effects.⁴

These important and innovative studies provide strong evidence that social spillovers matter. They do not attempt to estimate treatment-status-specific peer effects for all subjects using individual data. Our work leverages a setting in which spillovers may be estimated both for treated subjects and controls. We also analyze individual data for all subjects in order to explore the peer mechanism directly and extensively.

3 Experimental Design

The experiment was conducted in two waves. Wave I of the experiment entailed surveying 839 freshman students in Santa Catalina Residence Hall at the University of California, Santa Barbara. Students were recruited at a table outside the dormitory dining hall from January 14 to January 20 in 2011 and were paid \$5 for filling out a survey and giving consent to be contacted later. The survey included questions on demographics and background.⁵ To thank them for participating, students were entered in a lottery for an iPod Nano, retail value \$149. The dormitory had 1,178 residents, so the 839 students represented 71 percent of the student population in the residence hall. Potential selection issues are discussed in Section 4.2.

Wave II took place from January 27-February 3. 658 of the 839 Wave I students chose to participate in Wave II. During this wave, we invited all Wave I participants to

⁴ Breza and Chandrasekhar (2018) use the same network from Banerjee et al. 2013 and implement a savings monitors experiment to study a specific peer mechanism: individuals wanting to impress others through their actions. Similar to our approach, Dieye et al. (2018) use network data combined with experimental variation in cash transfers to study peer effects in school decisions in Colombia.

⁵ The survey and consent form are in Appendix 1.

complete an online friendship survey. The first question on the survey instructed students to click on the checkboxes next to the names of their friends. The alphabetized list included the names of all Wave I participants.⁶ They were shown at most 50 names per screen.⁷ The second question asked students to report their first best friend from the list of students they had just identified as friends. The third and fourth questions similarly asked them to identify their second and third best friends from their list of identified friends. They were given the option to report “none” for each of the three best friend questions.⁸ Finally, at the end of the survey we asked students to report their dormitory room number so that we could determine the treatment status of their roommate(s).

After subjects filled out the friendship survey they were instructed to return to the tables at the residence hall where we paid them \$8 for completing the friendship survey and presented them their treatment or control assignment. Subjects randomized into the treatment group were informed they would be paid \$60 if they visited the University Recreation Center (“Rec Center”) eight times or more between February 6 and March 6. They were also told that only one Rec Center visit per day would count. At the end of the experiment period, we paid subjects in the treatment group who attained the eight visit threshold \$60.

⁶ We chose to present them in alphabetical order rather than a random ordering so that students could easily find their friends on the list.

⁷ The initial page of the friendship survey read as follows: “Below are the names of some students in Santa Catalina you may or may not know, listed in alphabetical order by last name. We will show about 50 names at a time. Please make check marks next to the names of people you would say are your friends”. There were 17 pages of names. Each of the following pages after the initial page read as follows: “Here is the next page of names. Please make check marks next to the names of people you would say are your friends.”

⁸ The first best friend question read as follows: “Please select your best friend from the names you chose on the previous pages. If you chose no friends, select None.” The second best friend question read as follows: “Please select your second best friend from the names you chose on the previous pages. If you chose fewer than two friends, select None.” The third best friend question followed and was worded similarly.

Several unique aspects of the research design in this pay-for-exercise experiment merit comment, as they relate directly to the identification strategy. The recreation center is a 20 minute walk from the residence hall and a 10 minute walk from most campus lecture halls. Thus, a visit to the Rec Center involved a nontrivial commitment of time and effort. Subjects were recruited from the same dormitory so that they would be likely to know some of the other subjects in the experiment. More than 50 percent of the students living in the dormitory participated in the experiment. Students were surveyed in the middle of the academic year (January), so that having lived together in the same residence hall for several months, they would have had time to form meaningful friendships. Friends were elicited before students learned their treatment assignment. Thus, the reported friends were not influenced by treatment or treatment status. Random assignment of treatment served two purposes in this experiment. It induced random variation in exposure to treatment, which allowed identification of individual effects, and it induced random variation in exposure to treated and untreated peers, which allowed identification of spillover effects. We emphasize that the main purpose of the experiment was not to perform a full-scale evaluation of pay-for-exercise as a policy to improve long-run health outcomes on college campuses. The important issues of habit formation and health effects in interventions of this kind have already been explored carefully in previous work.⁹ The analysis here focuses on peer effects and the fundamental behaviors that drive them in an effort-elicitation setting.

4 Data

⁹ Charness and Gneezy (2009) find evidence of habit formation and health benefits in a short-term exercise intervention very similar to ours. Acland and Levy (2010) also find some evidence of habit formation, though the acquired habits fade quickly over time.

To construct the analysis sample we merge experiment data (Wave I demographic survey and Wave II friendship survey) and subject treatment status with gym visit information. Data on student visits to the Rec Center were obtained from the Rec Center database. These data include the time, date, and student identification card barcode of every gym entry. A valid student identification card (that includes a photo) is required for entry into the Rec Center. Students must present their card to the attendant at the front desk who electronically scans each card while the student simultaneously enters his unique seven digit student identification number into the key pad. The student photo is clearly visible to the attendant who scans the card and only one card can be presented per student. Thus, it was difficult for subjects in our experiment to cheat by using someone else's card or by scanning multiple cards for a single visit.¹⁰

4.1 Defining Peers

Not surprisingly, the Santa Catalina Residence Hall has a high degree of social connectedness. The average number of friends reported on the friend survey was 18. As such, every subject who reported having any friends is connected to every other subject by five or fewer degrees of separation (see Figure 1). In other words, students who were not direct friends with each other were nevertheless very likely to be friends of friends, or friends of friends of friends. In such an environment, collecting a comprehensive list of friends is not particularly helpful for defining meaningful friends, but it is helpful for focusing attention on friends to help subjects identify influential friends – first through

¹⁰ One might worry that students scanned their card multiple times in a day to accumulate visits. They were informed when they received their treatment status that only one visit per day would count. We also observe that in the pre-experiment period and experiment period, no subject exceeded one visit per day.

third best friends. Subjects designate their first, second and third best friends from their larger self-reported friendship network elicited in the friendship survey. Unless otherwise stated, we define “best friend” as the best friend stated by the subject. In the Facebook era – where it has been suggested that agents identify many people as friends who they do not interact with at all¹¹ – self-stated best friend is apt to be particularly meaningful, distinguishing a friend who has been integrated into the subject’s life from someone on the periphery.

That being said, we explore alternative peer definitions and the broader best friend network in Section 5.2. We consider several alternative measures and specifications. (1) Good friends: Second and third self-stated best friends. (2) Friend-stated best friend: The individual (or group of individuals) who report a subject as their best friend on the friend survey, but are not necessarily reciprocated relationships. (3) Reciprocated best friends: The relationships where both individuals report being best friends. (4) Roommates: The person (or two people) assigned to the subject’s dormitory room. (5) Lastly, we examine the overall fraction of friends treated. While this may be intuitively appealing, with a 60 percent treatment rate and a large number of reported friends (18 on average), there is very little variation in the fraction of friends treated.

4.2 Sample Characteristics

The final analysis sample includes 612 Santa Catalina residents because 46 of the 658 Wave II participants did not report a best friend and thus are not included in the analysis

¹¹ The Economist, February 26, 2009.

sample.¹² As such, 78 percent of the Wave I sample participated in Wave II (658 of 839 students) and 73 percent of Wave I (612 of 839 students) are included in the sample we analyze (selection issues are discussed below). Table 1 reports the summary statistics. Column 1 lists mean observable characteristics for the analysis sample. On average subjects report exercising 3.5 times per week, though this could include non-gym forms of exercise. From the Rec Center data, we see that in the month leading up to the experiment subjects made 2.2 visits on average.¹³ Sixty-five percent of subjects did not go to the Rec Center at all in the pre-experiment period.

In general, selection into experiment participation does not threaten the internal validity of the results since treatment status was randomly assigned; however, in column 2 we provide results from a series of balance regressions to assuage the concern that the 612 students in the final sample differed on observables from the 227 who filled out an initial survey but did not go on to fully participate in Wave II. There were no statistically significant differences between participants and non-participants, except that more males chose not to return for Wave II.

We similarly test for balance across treatment status to ensure that our treatment assignment was indeed random. Note that the randomization was designed to treat 60 percent of Wave I participants. Column 3 reports mean characteristics of observables for

¹² Even though the friendship survey instructed students to report a best friend based on the list of participants, best friend non-reporting could result from social isolation or because their best friend did not participate in the experiment; perhaps they were in another dorm, a non-participant in Santa Catalina, or attended another school. There were 12 students who did not report a first best friend but who reported a second or third best friend. In these cases we replaced the missing first best friend with the reported second best friend, and the third best friend if both the first and second best friend were missing. This adjustment will attenuate the estimated peer effect since a second or third best friend likely is less relevant than a first best friend. Our results are robust to the exclusion of these 12 students.

¹³ Although we observe Rec Center visits for the full month prior to the experiment period, we instead construct pre-experiment period month visits using the two weeks prior, which we then scale up to four weeks, to avoid non-representative gym behavior that may be associated with the first of the calendar year or the first few weeks of winter quarter.

the control group. Column 4 shows that the composition of the control and treatment groups is very similar; there are no statistically significant observable differences between control and treatment groups.

Despite the balance discussed above, one might also wonder if participation in Wave II was correlated with a best friend's treatment status or the treatment status of a roommate. Table 2 provides evidence that this is not a concern. Panel A (column 1) reports a 73 percent Wave II participation rate for Wave I participants with no treated roommates. Column 2 shows there is no difference in the participation rate between those with and without a treated roommate. To obtain this estimate, we regress an indicator for being in the analysis sample on an indicator for having a treated roommate. Panel B repeats this exercise replacing the treated roommate indicator with an indicator for having at least one treated friend-stated best friend. Again, that treatment status of a friend-stated best friend does not predict participation in Wave II. Note that Panel B uses friend-stated best friends and not self-stated best friends because we only observe self-stated best friends for subjects who participate in Wave II. Finally, one might worry that students whose "real" best friend was not in the experiment reported an alternative best friend who was on the list but is less of a relevant peer. This behavior will attenuate our estimated peer effects, thus any effects we detect will be understated. This is similarly true for individuals whose true best friend does not live in Santa Catalina.

5 Results

5.1 Individual effects

Before we investigate whether the intervention caused effort spillovers, we first document the overall effect of treatment on effort outcomes with a simple comparison of means. Throughout the analysis, we focus on two outcomes: The first is the number of times the subject visited the Rec Center in the treatment period. Figure 2a shows Rec Center usage for treated and control subjects before and during the treatment period. Control subjects and treated subjects made an average of 2.3 and 2.4 recreation center visits, respectively, in the 4-week pre-experiment period. During the 4-week treatment period, control subjects visited the recreation center an average of 2.4 times, while treated subjects increased their usage dramatically to 7.3 mean visits. This 4.7 visit difference in differences is statistically significant at conventional levels.¹⁴

The second effort measure is whether the subject visited the Rec Center at least 8 times in a 4-week period, as would be required for a treated subject to receive payment in the treatment period. Figure 2b shows the fraction of subjects who reached the 8-visit threshold for treated and control subjects before and during the treatment period. 11 percent of control subjects and 12 percent of treated subjects went to the Rec Center 8 times or more in the 4-week pre-experiment period. During the 4-week treatment period, the fraction of control subjects visiting the recreation center 8 times or more stayed constant at 10 percent, whereas fully 62 percent of treated subjects reached the payment threshold. Again, this difference is significant at conventional levels. We conclude that treatment incentivized subjects to significantly increase their Rec Center visitation rate.

5.2 Spillover Effects

¹⁴ To assuage concerns that students may have been going to the gym on their way home from class, a potentially less costly gym visit compared to one that is initiated from Santa Catalina Residence Hall, we restrict the sample to weekends and find a similar treatment effect (see Appendix Figure 1).

Next, we estimate both the direct effect of the treatment and any spillovers that may exist.

Our main results come from the following regression:

$$y_{it} = \alpha_1 + \beta_1 T_{it} + \beta_2 T_{it}^{BF} + \beta_3 T_{it} \times T_{it}^{BF} + \gamma_i + \delta_t + \epsilon_{it} \quad (1)$$

where y_{it} is the outcome for subject i in period t . T is an indicator equal to one if the subject is treated and T^{BF} is an indicator equal to one if the subject's best friend is treated. γ_i and δ_t are individual and period fixed effects, respectively. The inclusion of individual fixed effects controls for all time-invariant observable and unobservable characteristics of a subject, and the period fixed effect controls for temporal changes in gym behavior over the two periods. ϵ_{it} is the usual idiosyncratic error term. All standard errors are clustered at the individual level.

β_1 measures the direct effect of treatment; that is, the treatment effect for a subject with a control best friend. β_2 captures the spillover effect between treated and control subjects, or the effect of a treated best friend for a control subject (relative to a control subject with a control best friend). The interaction term β_3 represents the spillover effect between treated individuals, it is the additional effect of a treated best friend for a treated subject (relative to a treated subject with a control best friend). Thus, $\beta_2 + \beta_3$ is the total effect of a treated best friend on a treated subject (relative to a treated subject with a control best friend). The omitted group is control subjects with a control best friend.

Table 3 reports estimates for Equation 1. Column 1 shows that students who were treated with a control best friend went to the gym 4.0 times more than their untreated counterpart. This estimated direct effect of the treatment is similar to the simple

differences in means effect reported in Section 5.1 (4.7 visits). Interestingly, we find no statistically significant differences in Rec Center visits between control subjects with treated best friends and control subjects with control best friends; β_2 is small and indistinguishable from zero.¹⁵ In contrast, there is an economically and statistically significant spillover among the treated. Treated subjects made 1.0 more visit to the Rec Center if their best friend was also treated. Both the interaction (spillover) and the total effect are statistically significant at conventional levels.

Results for the threshold outcome (column 2) follow a similar pattern. Treated subjects were 7.1 percentage points more likely to reach the 8-visit threshold if their best friend was treated, but this estimate is imprecise. The imprecision of the threshold treatment spillover point estimate can more easily be seen in Figure 3. This figure plots the cumulative distribution function (CDF) of the difference in gym visits (experiment period – pre-experiment period) for each of the four groups. A Mann-Whitney test rejects the null hypothesis that the samples of treated subjects with treated and control best friends are from the same population with a p-value < 0.01, the CDF for treated subjects with treated best friends is slightly to the right of the CDF for treated subjects with control best friends.¹⁶ We probe this result further in Section 5.5 and show that for the types of students most likely to be sensitive to the treatment incentive, having a treated best friend significantly increased the probability of reaching the payment threshold.

Next we ask whether one's first best friend has special influence that other close friends do not have, or if other good friends produce similar spillovers. While discussing this specification we will refer to the first reported best friend as the 'best friend' and the

¹⁵ While we do not have information about how the incentive payments are spent, this result is consistent with treated friends not sharing their incentive with untreated friends to encourage gym going.

¹⁶ A person with 2 visits in the pre-experiment period would need a 6 visit increase to reach the threshold.

second and third reported best friends as ‘good friends’. To explore this we take the subsample of subjects with at least one treated best friend or good friend. This facilitates comparison between subjects with treated best friends to those with only a treated good friend, to those with both treated best and good friends. This restriction only eliminates 55 subjects; almost everyone has at least one treated best or good friend. The omitted category in Table 4 is control subjects with a treated good friend but not a treated best friend. To facilitate comparison with Table 3, column 1 in Table 4 replicates column 1 from Table 3. Not surprisingly, given the similar samples, the results in column 1 are essentially the same as those reported in column 1 of Table 3.¹⁷ Column 2 of Table 4 drills down a level by including an additional indicator for subjects with both a treated best friend and a treated good friend. This allows us to ask whether having a treated good friend in addition to the best friend differentially changes gym attendance.¹⁸

We find no evidence that having a treated good friend in addition a treated best friend differentially increases gym attendance. The bottom panel of Table 4 focusses on treated subjects (control comparisons are omitted to conserve space since they are never distinguishable from zero) and reports comparisons between various treatments. Following from previous comparisons, a treated subject with a treated best friend visits the gym 1.4 more times than a treated subject with only a treated good friend. This comparison is statistically indistinguishable from the additional visits by a treated subject with a treated best and good friend relative to a treated subject with only a treated good

¹⁷ The only difference between these two columns is that in Table 3 the omitted group includes those with no treated best or good friends, while those subjects are excluded from the sample in Table 4.

¹⁸ One could alternatively ask whether increasing the fraction of friends treated increases gym attendance for those who are treated and untreated with a best friend who is treated or untreated. With the mean number of reported friends being 18, there is little variation in the fraction of friends treated. We therefore prefer the best and good friends specification reported in Table 4, but report the fraction of friends specification in Appendix Table 1.

friend. In contrast, the difference in average gym visits for treated subjects with a treated best friend and good friend relative to and those with only a treated best friend is small in magnitude and indistinguishable from zero. These results are consistent with the first best friend being pivotal.

While we have thus far focused on self-stated best friends, the most relevant peer remains an open question in this literature. It could be that proximity matters in ways that mean that subjects are influenced by roommates. It may also be that directionality of friendship ties matter. Is the best friendship reciprocated, or does it only go one way? There is evidence from psychology showing that reciprocated friendships are more intimate and, as such, produce more effective peer-influence than unilateral ties (Almaatouq et al., 2016). To investigate these possibilities, we partition the definition of best friend to indicate whether or not the self-stated best friend can also be categorized under an alternative peer category. More specifically, we define an alternate peer two ways: in the first specification it is defined as roommates and in the second specification, friend-reported best friends. The specifications take the following general form:

$$y_{it} = \alpha_1 + \beta_1 T_{it} + \beta_2 T_{it}^{BF} + \beta_3 T_{it}^{AP} + \beta_4 T_{it}^{BF_APS} + \beta_5 T_{it}^{BF_APD} + \beta_6 T_{it} \times T_{it}^{BF} \quad (2)$$

$$+ \beta_7 T_{it} \times T_{it}^{AP} + \beta_8 T_{it} \times T_{it}^{BF_APS} + \beta_9 T_{it} \times T_{it}^{BF_APD} + \gamma_i + \delta_t + \epsilon_{it}$$

where *BF* denotes a best friend who is not the alternate peer (best friend but not roommate), *AP* denotes alternate peer is not the best friend (roommate but not best friend), *BF_APS* denotes best friend is also an alternate peer (best friend is a roommate), and *BF_APD* denotes best friend and an alternative peer where they are not the same person (best friend and a roommate). Note that these are mutually exclusive groups. The

omitted group is a control subject with no treated alternative peers with an untreated self-stated best friend.

In the case of roommates, the alternative peer is considered treated if the subject has at least one treated roommate and is untreated if the subject has no treated roommates, either because the roommate(s) is randomly assigned the control group or because she did not participate in the experiment.¹⁹ In the case of friend-stated best friends, subjects could have zero friend-stated best friends if no one claimed them as a best friend (40% of subjects) or up to four (3 of the 612 participants had 4 other participants claim them as best friends). As such, in the second specification the alternative peer is coded as treated if the subject has at least one treated friend-stated best friend and is coded as untreated if a subject's friend-stated best friend(s) is either assigned the control or if no one claims them as a best friend.²⁰ Finally, T^{BF_APS} equals one if a subject has a treated alternative peer who is also a self-stated best friend (i.e., they are the same person), and T^{BF_APD} equals one if a subject has a treated alternative peer and a treated self-stated best friend, but who are not the same individuals. Table 5 reports the fraction of subjects in each of these categories.

The results for specification (2) are reported in Table 6. Column 1 reports the results when the alternative peer is defined as a roommate and column 2 reports the results when the alternative peer is defined as a friend-reported best friend. While we do not have sufficient statistical power, given our substantial sub-grouping, to reject the

¹⁹ Typically, there are two students per dorm room but 6% of students share a room with two others. Also, most roommates in the Santa Catalina Residence Hall are randomly assigned, conditional on a few requested characteristics.

²⁰ Lack of reciprocation is a common feature of social networks and not an idiosyncrasy of the Santa Catalina Residence Hall. In the National Longitudinal Survey of Adolescent Health, for example, best friend nominations are reciprocated about a third of the time, as here (Card and Giuliano, 2011).

equality of treatment effects across all groups at conventional levels, the spillover is generally biggest for those with both a treated self-stated best friend and a treated alternate peer. To be precise, we cannot reject that the three categories with treated self-stated best friends differ in gym visits from each other in either column. On the other hand, there is somewhat noisy evidence that having both a treated self-stated best friend and a treated alternate peer leads to more gym visits than just having a treated alternate peer (p-values for these comparisons range from 0.03 to 0.18). We view this auxiliary exercise as rendering estimates that are consistent with the idea that spillovers between treated subjects are sensitive to the closeness of relationships.

A first pass at the data produces several main findings: (1) Subjects go to the Rec Center 4 more times in the treatment period if treated. (2) Treated subjects visit the Rec Center 1 more time if their best friend was randomly assigned to treatment. (3) A subject's best friend, particularly those who are close best friends (when the best friend is also a roommate or when the best friendship is reciprocal), likely matter most. (4) There is no visible peer effect for control subjects. It is worth noting that the observed peer effect, a spillover from treated subjects to treated subjects, would not have been identified by a standard partial population experimental design, where only spillovers from treated subject to controls are identified.²¹

5.3 The Influence of Best Friend Behavior on Own Behavior

²¹ To explore spillovers, we have focused on a fundamental and objectively measurable behavioral response: observed effort outcomes. We do not make strong claims about whether increased Rec Center visits led to better long-run health outcomes or led to more exercise, overall. It is possible that students substituted from one form of exercise to another, or that they came to the Rec Center and did little or no exercise. Going to the Rec Center, itself, required effort.

Peer studies often attempt to answer the question: How much does peer behavior influence own behavior? In our setting, the task is to estimate the effect of best friend Rec Center visits on own visits. However, peer selection is likely a serious confounder. Students who frequent the Rec Center may choose best friends who do the same.

Consider a simple model that elucidates these challenges. Let $i=1$ or 2 be the subject and the subject's best friend, respectively. Let y_{it} be i 's effort outcome in period t , T_i be treatment status, and ε_{it} the usual error term. Suppose the true causal structure is:

$$y_{1t} = \alpha_1 + \beta_1 y_{2t} + \delta_1 T_{2t} + \gamma_1 + \delta_t + \epsilon_{1t} \quad (3)$$

$$y_{2t} = \alpha_2 + \beta_2 y_{1t} + \delta_2 T_{1t} + \gamma_2 + \delta_t + \epsilon_{2t} \quad (4)$$

In the canonical simultaneous equation model, the exclusion rule for identification requires that there be at least one exogenous variable excluded from each equation. Usually, it is a challenge to find exogenous variables to satisfy the exclusion restrictions. Here, T_2 is excluded from equation (3) and T_1 is excluded from equation (4), and random assignment of T_1 and T_2 obviates the usual concerns about selection. Thus, the exclusion restrictions here may be more plausible than in peer studies that either lack random assignment or do not feature it at the individual level. The exclusion restriction is interpretable as the assumption that a subject's treatment assignment influences his own effort outcome but has no direct influence on his friend's outcome. It is possible, of course, that the exclusion restriction is violated.²² We do not insist that it is valid because we view the IV regressions as offering a convenient normalization/scaling by which to interpret the findings: They tell us the size of the spillover from best friend's treatment if we imagine that it functioned entirely through best friend's visits.

²² For example, the restriction is violated if individuals effectively encourage or exhort their best friends to go to the Rec Center more often, even when they themselves do not.

In the peer effect variant of the simultaneous equation model, coefficients in the two structural equations are often assumed to be identical.²³ We do not make that assumption here because subjects often denote best friends who do not reciprocate the bond. For instance, estimating β_1 requires defining the best friend as the self-stated best friend while estimating β_2 requires defining the best friend as friend-stated best friends. Standard IV techniques allow us to estimate this set of structural equations. As we are only interested in β_1 , we employ an IV approach to estimate Equation 2 only.

One other subtle deviation from the standard peer model merits comment. The evidence in Table 3 indicates that spillovers differ by treatment status. Thus, we wish to estimate a *treatment-specific* β_1 . To do this, we deviate slightly from the structural equation presented in Equation 3 and instead estimate an interacted model which includes interactions with subject's treatment status as in Equation 1.

The first and second stage results are reported in Table 7. The first stage results show that the best friend's treatment assignment (the instrument), predicts the endogenous regressor with a high degree of statistical significance, the F-statistic is 37. Panel B reports the second stage results: a treated subject's effort choice rises by 0.25 visits when his best friend's effort choice rises by 1 visit.

The IV regressions indicate that if one imagines the spillover as resulting entirely from best friend's visits, then the spillover is about 25 percent as large as the direct effect of the intervention. In Table 3, the results of the OLS regressions of individual visits directly on best friend's treatment assignment showed that subjects visited the Rec Center about 1 more time if their best friend was treated. This was about 25 percent of the overall effect of the intervention (a 4-visit increase in Rec Center usage). Interestingly,

²³ See, for example, Moffitt (2001).

the size of the spillover is similar, whether one imposes assumptions that force best friend's treatment assignment to work only through best friend's outcomes, or whether one allows it to have a direct effect, as well.

5.4 Mechanisms

What, then, is the mechanism for the peer effect? In the literature it is often difficult to discern the peer mechanism, even when a peer effect has been isolated or identified. In partial population experiments that capture mean responses to mean behavior of a large group of peers, investigating peer mechanisms that operate at the individual level has proved to be challenging. In the broader literature as well, peer influence has tended to be a black box phenomenon. Detailed data on Rec Center visits allow us to look more closely at mechanisms than is sometimes possible. We will focus on three broad categories of potential mechanisms.

1. Coordination. Friendships allow for coordination that may facilitate completion of an effort-intensive task. We highlight two ways in which coordination could be decisive. In models of self-control and pre-commitment, individuals fail to meet goals because the present self lacks the ability to bind the future self to a plan of action; the present self would rather engage in an activity that is more immediately pleasurable (O'Donoghue and Rabin, 1999, 2001). Having a treated best friend could remedy this problem. If one coordinates by committing oneself to exercise with a friend, it is more difficult for one's future self to back out. In short, individuals who have both been incentivized may use each other to devise commitment mechanisms.

A second variant of the coordination mechanism features complementarities in the production of utility. Rec Center visits with a friend may produce more utility (or reduce the effort costs of Rec Center usage) relative to single-person visits. This is the standard concept of complementarity in production (e.g., Lazear, 2000) applied to a context in which the produced good is utility (Stigler and Becker, 1977). If it is less onerous (or more fun) to go the gym with a best friend, then having a treated best friend (who visits the Rec Center more, *ceteris paribus*) provides more opportunities for lower cost visits.

2. Imitation. A broad class of models commonly used in empirical studies of peer effects posits that individuals seek to imitate the expected behavior of others in their reference group. In this framework, individuals seek to minimize the difference between their own effort choices and the effort choices of peers because they derive utility from imitation or sameness (Akerlof, 1997; Akerlof and Kranton, 2000). Imitation models have been taken seriously by researchers and policymakers across disciplines. Concerns about “peer pressure” in adolescent drinking, smoking, and obesity, for example, are often based (formally or informally) on the notion that students imitate others who exhibit these behaviors.

3. Information Exchange. A third broad class of models and empirical analyses of networks and social interactions emphasizes learning and information exchange (e.g., Duflo and Saez, 2003; Miguel and Kremer, 2007). Much work, particularly in the development literature, centers on diffusion across networks of information related to health or production technology. See Duflo (2006) for a survey of this literature. An information exchange story in our setting would be that best friends, when treated, communicate information to subjects that facilitates Rec Center visits.

We begin with coordination. In the Rec Center data, every visit has a time stamp. Figure 4 shows a plot of average Rec Center visits per person by day of visit for the 4 groups of subjects in Table 3 (based on treatment assignment and best friend treatment assignment). The zero point on the horizontal axis denotes the beginning of the treatment period. Main findings from Table 3 are visible in the figure. In the top row of graphs, which depicts outcomes for control subjects, the plots look very similar. For controls, best friend treatment assignment does not alter the number or timing of visits. In the bottom row of graphs, which depicts outcomes for the treated, the plots appear similar during the pre-experiment period but diverge in the treatment period. For treated subjects, there were more visits in the treatment period if the best friend was treated.

In Table 8, we tighten the focus on the timing of visits. We define a simultaneous visit for a subject as one in which the time-stamp for a subject's visit was within a few minutes – 1, 5, 10, 15 and 20 minutes – of a subject's best friend's visit. The table reveals that treated subjects with treated best friends made about 1.1 more simultaneous visits with their best friend during the treatment period if their best friend was treated (compared to treated subjects with control best friends). This result is statistically significant at conventional levels.

But it is possible that this difference in means does not indicate deliberate coordination of visits. Treated best friends go to the Rec Center more often, other things equal, and so the probability of a simultaneous visit is higher through pure chance when a subject's best friend is treated. We created placebo best friends, in order to investigate whether unplanned simultaneous visits could account for the mean difference estimated in Table 8. Each treated subject is randomly reassigned a best friend from all possible

treated subjects. We repeat this random reassignment 1,000 times and plot the distribution of the estimated relevant spillover effect along with the corresponding t-statistics. Figure 5 reports these results. Indeed, the number of simultaneous visits with a randomly assigned placebo best friend is about a tenth of a visit higher if the placebo best friend is treated. But this effect is clearly quite small.²⁴ Although the number of “accidental” simultaneous visits is slightly higher when a subject’s best friend is treated, this accounts for only a small fraction of the observed difference in simultaneous visits with true best friends. In brief, the evidence indicates that treated subjects visited the Rec Center jointly with a best friend about 1 additional time if the best friend is treated, and that unplanned coincidental visits do not account for this difference.

Table 3 showed that having a best friend assigned to treatment led to one more visit *overall*. The results in Table 8 imply that having a best friend assigned to treatment led to 1 more *joint visit*. Thus, simultaneous visits can account for the entire observed peer effect on visits. While this result is consistent with coordination, it seems to argue against imitation because subjects who derive utility strictly from imitation need not visit the Rec Center at the same time in order to imitate peer behavior. The “contagiousness” of this behavior appears driven entirely by joint visits, by activities agents engage in together, rather than through utility derived from identity or self-definition, as posited in imitation models. Moreover, if imitation was the main channel, one might expect a spillover between control subjects and treated best friends, and we do not detect such an effect. Further, there would seem to be little room for the information story. Very little

²⁴ For the placebo exercise, we use a ten minute window but the results are robust to using a 1, 5, 15 or 20 minute window as well.

information was needed to visit the Rec Center: in the experiment, subjects were informed of the location of the Rec Center, and the vast majority of them already knew.

Overall, the evidence appears most consistent with coordination rather than imitation or information. But it is important to note that coordination can arise either because of commitment problems or complementarities in utility. To the extent that commitment difficulties are the driving factor, the evidence of coordination may also help explain why there is no observed spillover from treated subjects to control best friends. For control subjects (who would then have no intrinsic desire to go to the Rec Center more often), there would exist no commitment problem or other obstacle to be overcome through coordination, no conflict between present and future self, and thus *no marginal gain to having a treated best friend*. Treated subjects, in contrast, would desire to go to the Rec Center precisely because of the financial incentive provided, and would thus benefit from having a best friend to help them muster the discipline to do it.

5.5 Heterogeneity

Finally, one might wonder if certain types of individuals are more responsive to interventions of this nature. We explore this possibility by measuring heterogeneous treatment effects and spillovers by previous Rec Center usage. Figure 2 shows that subjects visited the Rec Center 2.3 times on average during the pre-experiment period. However, a large fraction of subjects (65%) did not visit the Rec Center at all in the month before the treatment period. Conditional on going at least once, subjects averaged 5 pre-experiment period Rec Center visits. These two groups, who we call “Users” and “Non-users”, may be expected to respond in different ways to best friend treatment

status. Specifically, subjects without a history or habit of Rec Center usage may be more dependent on external motivators to get them there than subjects who used the Rec Center in the absence of financial incentives. In short, Non-users may be the marginal subjects, those for whom a peer nudge could be decisive.

Table 9 replicates Table 3, except that we divide subjects into four groups, based on their pre-experiment Rec Center usage: Rec Center Non-users whose best friends are also Non-users; Non-users whose best friends are Users; Users with Non-user best friends; and Users with User best friends. The results reveal that treated Non-users with Non-user best friends visited the Rec Center about 2 more times and were 22 percentage points more likely to reach the payment threshold when the best friend was treated. Both results are statistically significant at conventional levels. No other pairings produced spillovers large enough to be statistically discernable at conventional levels. Following Dobbie and Fryer (2015), we report corrected p-values for testing multiple hypotheses using the Holm step-down method outlined in Romano, Shaikh and Wolf (2010), and obtain nearly identical results. This is a conservative correction. As such, we have a high degree of confidence in estimates that remain statistically significant after the correction.

6 Summary

Strong claims have been made about the “contagiousness” of health-related behaviors and outcomes, but truly exogenous variation in behaviors has been hard to come by. A major goal of our field experiment has been to bring exogenous variation in behaviors to an endogenous set of connections and thus to learn about spillovers in effort elicitation settings and the mechanisms that drive them. We elicited best friends among college

students and offered monetary incentives for using the Rec Center to a treated subset. We found that treated students with treated best friends increased their use of the Rec Center more than treated students with control best friends. Control subjects did not alter their recreation center use at all, and were not influenced by the treatment status of their best friend. The effect was largest for students likely to be on the margin, those without a pre-existing habit of Rec Center usage. Finally, we provide evidence that the results are likely the result of coordination among students: Treated subjects made more joint visits with treated best friends.

The setting is stylized and was designed to investigate fundamental behavioral patterns. We have not offered a full-scale analysis on the effectiveness of short-term exercise interventions in changing long-run habits or health outcomes, as these issues have been explored carefully in previous work. However, in a more general way, the findings here may offer insights for how to improve the effectiveness of a variety of targeted interventions that seek to elicit effort from people in social settings. The large peer effect between the treated, particularly among subjects on the margin, may suggest that interventions designed to alter behaviors or elicit effort would be more cost-effective if they targeted and saturated a small number of networked populations than if they targeted small portions of many different networked populations. In the former approach, there would be a higher incidence of strong social ties between treated individuals, and this could generate significant spillovers that would otherwise not occur. Some examples of interventions that feature effort-elicitation in a social context (and about which this insight may be relevant) include: programs that try to induce people to lose weight, to raise test scores, to increase college enrollment, to donate blood, to get vaccinated, to use

contraceptives, to find jobs, and to improve sanitation practices. However, findings also suggest that even when peer effects amplify the effectiveness of an intervention among the treated, they may fail to spread to those who were not treated.

References

- Acland, D. and Levy, M. (2010), "Habit Formation and Naivete in Gym Attendance: Evidence from a Field Experiment," mimeo, UC Berkeley.
- Akerlof, G. (1997), "Social Distance and Social Decisions," *Econometrica*, 65(5): 1005-1027.
- Akerlof, G., and Kranton, R. (2000), "Economics and Identity," *Quarterly Journal of Economics*, 115(3): 715-753.
- Almaatouq, A., Radaelli, L., Pentland, A. and Shmueli, E. (2016), "The Role of Reciprocity and Directionality of Friendship Ties in Promoting Behavioral Change," *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, 33-41.
- Angelucci, M., and De Giorgi, G. (2009), "Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles' Consumption?" *American Economic Review*, 99(1): 486-508.
- Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. Jackson (2013): "The Diffusion of Microfinance," *Science*, 341 (6144).
- Bhattacharya, D., Dupas, P., Kanaya, S. (2018), "Demand and Welfare Analysis in Discrete Choice Models under Social Interactions" Working Paper.
- Bobonis, G. J., and Finan, F. (2009), "Neighborhood Peer Effects in Secondary School Enrollment Decisions", *The Review of Economics and Statistics*, 91(4): 695–716.
- Boisjoly, J., Duncan, G., Kremer, M., Levy, D., and Eccles, J. (2006), "Empathy or Antipathy? The Impact of Diversity," *American Economic Review*, 96(5): 1890–1905.
- Breza, E. and Chandrasekhar, A., (2018), "Social Networks, Reputation, and Commitment: Evidence from a Savings Monitors Experiment," Working Paper.
- Carrell, S. , Hoekstra, M., and West, J. (2011) "Is poor fitness contagious?: Evidence from randomly assigned friends," *Journal of Public Economics*, 95(7-8):, 657-663
- Card, D. and Giuliano, L. (2011), Peer Effects and Multiple Equilibria in the Risky Behavior Of Friends," NBER Working Paper 17088.
- Carrell, S., Fullerton, R., and West, J., (2009), "Does Your Cohort Matter? Estimating Peer Effects in College Achievement," *Journal of Labor Economics*, 27(3), 439-464.
- Charness, G., and Gneezy, U. (2009), "Incentives to Exercise" *Econometrica*, 77(3): 909–931.
- Christakis, N. and Fowler, J., (2007), "The spread of obesity in a large social network over 32 years," *New England Journal of Medicine*, 357(4): 370-9.
- Christakis, N. and Fowler, J., (2008), "Estimating Peer Effects on Health in Social Networks," *Journal of Health Economics*, 27(5): 1386-1391.

- Cohen-Cole, E. and Fletcher, J. (2008a), “Detecting Implausible Social Network Effects in Acne, Height, and Headaches: Longitudinal Analysis. *British Medical Journal*, 337: a2533.
- Cohen-Cole, E. and Fletcher, J. (2008b), “Is Obesity Contagious? Social Networks vs. Environmental Factors in the Obesity Epidemic,” *Journal of Health Economics*, 27 (5): 1382-1387.
- Dieye, R., Djebbari, H., and Barrera-Osorio, F. (2018), “Accounting for Peer Effects in Treatment Response,” Working Paper.
- Dobbie, W. and Roland, F. (2015), “The Medium-Term Impacts of High-Achieving Charter Schools” *Journal of Political Economy*, 123(5): 985-1037.
- Duflo, E. (2006), “Field Experiments in Development Economics,” Discussion paper, World Congress of the Econometric Society.
- Duflo, E., and Saez, E. (2003), “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment,” *Quarterly Journal of Economics*, 118(3), 815–842.
- Fischer, Stefanie J. (2017), “The Downside of Good Peers: How Classroom Composition Differentially Affects Men’s and Women’s STEM Persistence,” *Labour Economics*. 46: 211-226.
- Foster, G. (2006), “It’s not your peers, and it’s not your friends: some progress towards understanding educational peer effects,” *Journal of Public Economics*, 90 (8-9): 1455-1475.
- Hoekstra, M., Mouganie, P., and Wang, Y. (2018), “Peer quality and the academic benefits to attending better schools,” *Journal of Labor Economics*, 36(4): 841-884.
- Kremer, M. and Levy, D. (2008), “Peer Effects and Alcohol Use Among College Students,” *Journal of Economic Perspectives*, 23(3): 189-206.
- Kuhn, P., Kooreman, P., Soetevent, A.R., and Kapteyn, A. (2011), “The Effects of Lottery Prizes on Winners and their Neighbors: Evidence from the Dutch Postcode Lottery,” *American Economic Review*, 101(5): 2226-2247.
- Lalive, R. and Cattaneo, A. (2009), “Social Interactions and Schooling Decisions,” *The Review of Economics and Statistics*, 91(3): 457-477.
- Lazear, E. (2000), “Personnel Economics and Economic Approaches to Incentives,” *HKCER Letters*, 61: 1-8.
- Leider, S., Möbius, M., Rosenblat, T. and Do, Q. (2009), “Directed Altruism and Enforced Reciprocity in Social Networks,” *The Quarterly Journal of Economics*, 124 (4): 1815-1851.
- Lyle, D. (2007), “Estimating and Interpreting Peer and Role Model Effects from Randomly Assigned Social Groups at West Point,” *The Review of Economics and Statistics*, 89 (2): 289–99.

Miguel, E. and Kremer, M. (2004), “Worms: identifying impacts on education and health in the presence of treatment externalities.” *Econometrica* 72(1): 159-217.

Miguel, E. and Kremer, M. (2007), “The Illusion of Sustainability,” *Quarterly Journal of Economics*, 122(3): 1007-1065

Moffitt, R. A. (2001), “Policy Interventions, Low-Level Equilibria, and Social Interactions,” in *Social Dynamics*, ed. by S. N. Durlauf, and H. P. Young, pp. 45–82. MIT Press.

O’Donoghue, T. and Rabin, M. (1999), “Doing it Now or Later,” *American Economic Review*, 89(1): 103-124.

O’Donoghue, T. and Rabin, M. (2001), “Choice and Procrastination,” *Quarterly Journal of Economics*, 116(1): 121-160.

Oster, E. and Thornton, R. (2012), “Determinants of Technology Adoption: Private Value and Peer Effects in Menstrual Cup Take-Up,” *Journal of the European Economic Association*, 10(6):1263–1293.

Philipson, T. (2000), “External Treatment Effects and Program Implementation Bias,” NBER Technical Working Paper No. 250.

Romano, J., Shaikh, A., and Wolf, M. (2010), “Hypothesis Testing in Econometrics,” *Annual Review of Economics*, 2(1):75-104.

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

Stigler, G., and Becker, G. (1977), “De Gustibus Non Est Disputandum,” *American Economic Review*, 67(2): 76-90.

Yakusheva, O. Kapinos, K., and Eisenberg, D. (2014), “Estimating Heterogeneous and Hierarchical Peer Effects on Body Weight Using Roommate Assignments as a Natural Experiment,” *Journal of Human Resources*, 49(1): 234-261.

Zimmerman, D. (2003), “Peer Effects in Academic Outcomes: Evidence from a Natural Experiment,” *The Review of Economics and Statistics*, 85(1): 9–23.

Figure 1. Friendship Network
Degrees of Separation (Relative Frequency), Mean = 2.9

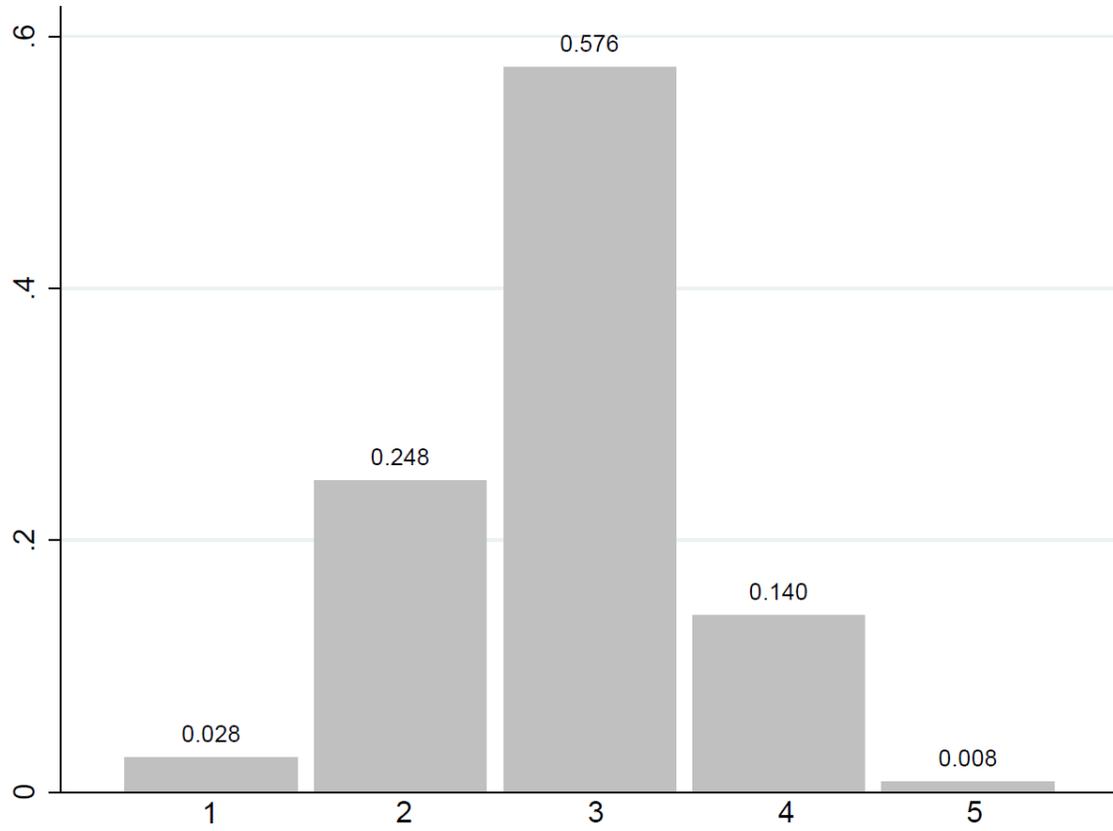
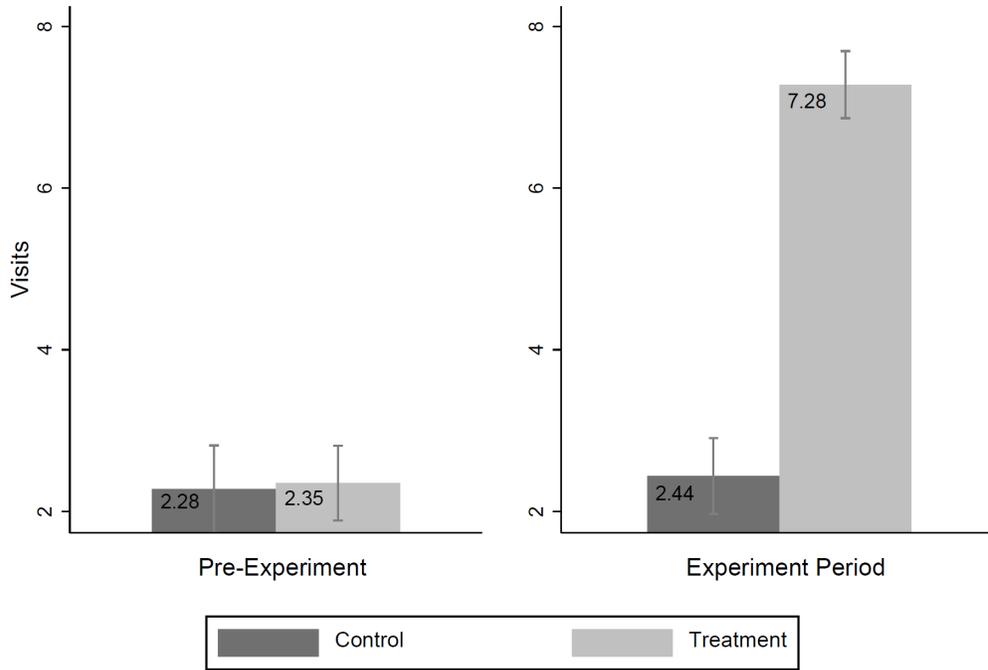


Figure 2.

A. Rec Center Visits (per/month)



B. Payment Threshold (>7 visits in the month)

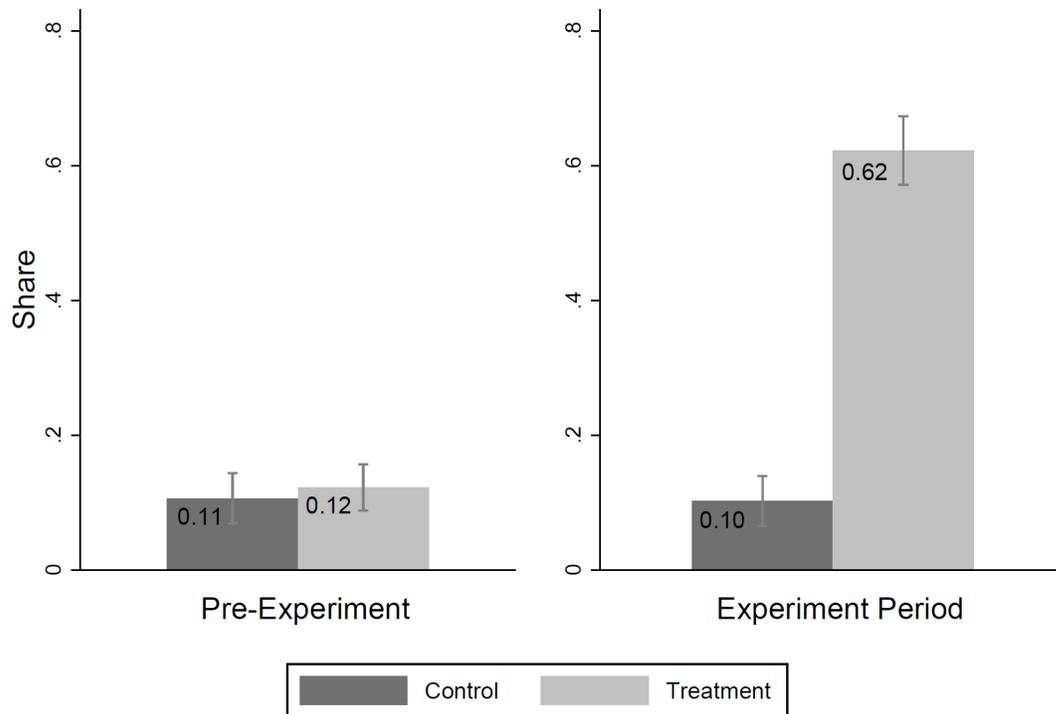
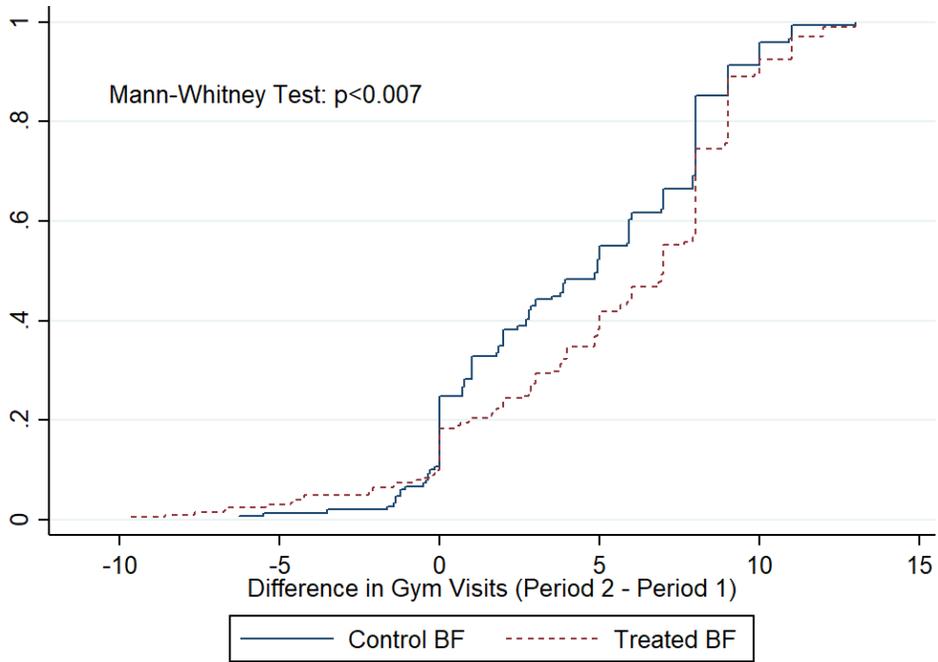


Figure 3.

CDF for Treated Subjects



CDF for Control Subjects

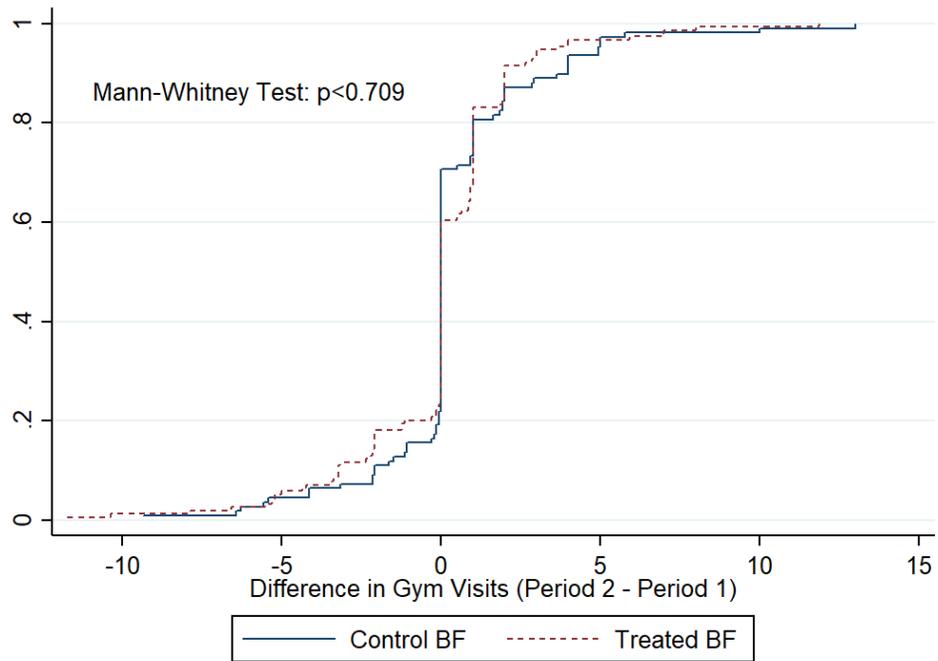


Figure 4. Average Rec Center Visits per Person by Day of Visit

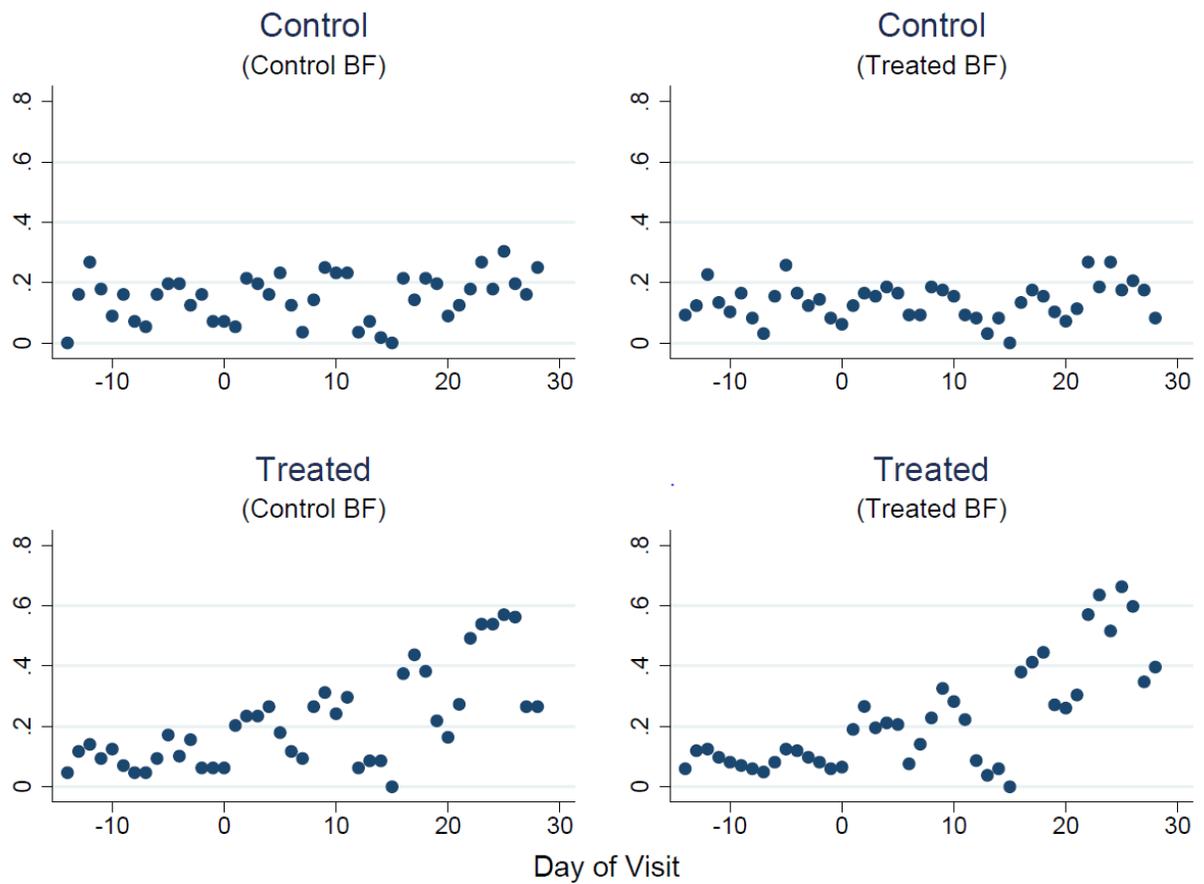


Figure 5. Placebo Test – Number of Simultaneous Visits with Randomly Assigned Best Friend

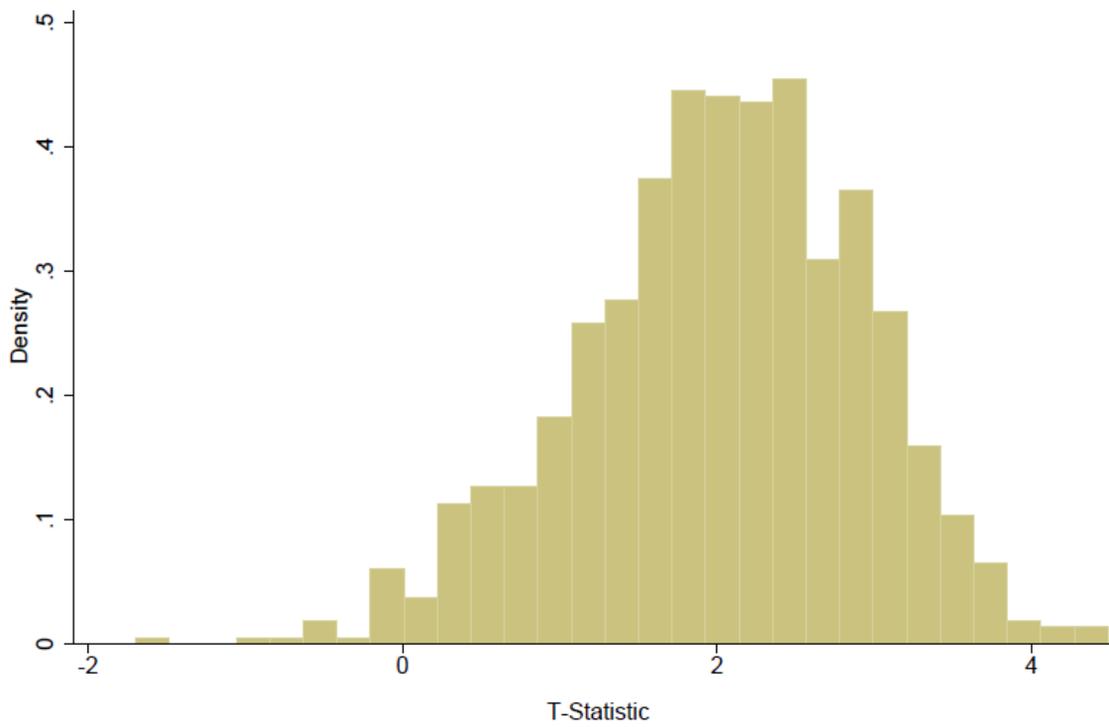
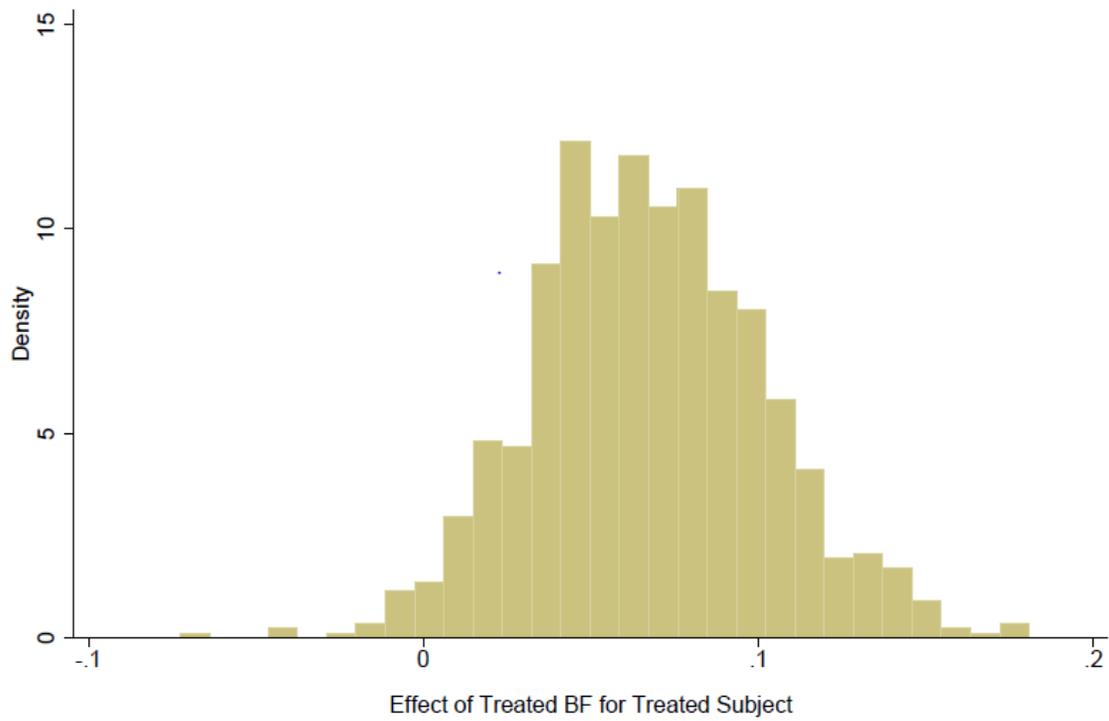
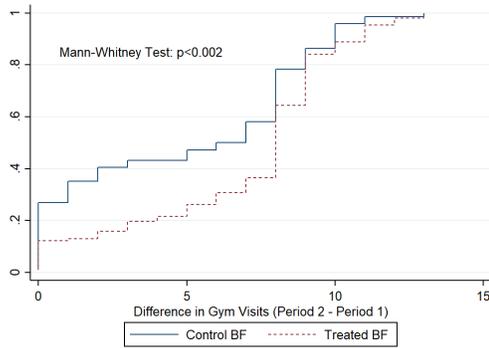
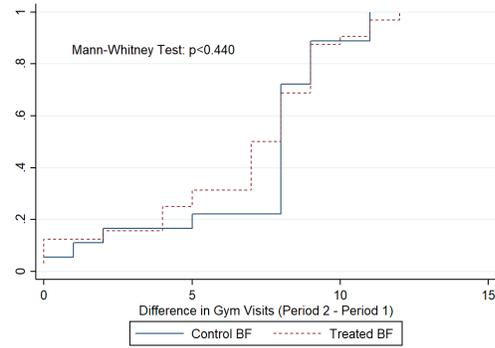


Figure 6. CDF for Treated Subjects

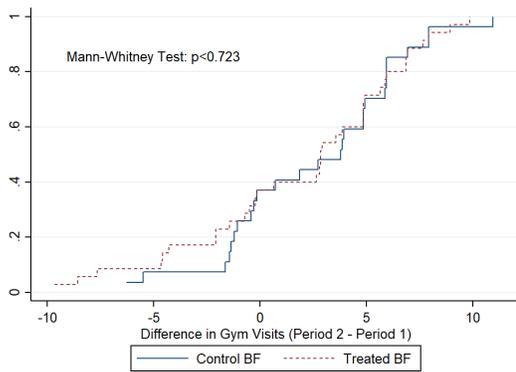
Subject: No Visits, BF: No Visits



Subject: No Visits, BF: Some Visits



Subject: Some Visits, BF: No Visits



Subject: Some Visits, BF: Some Visits

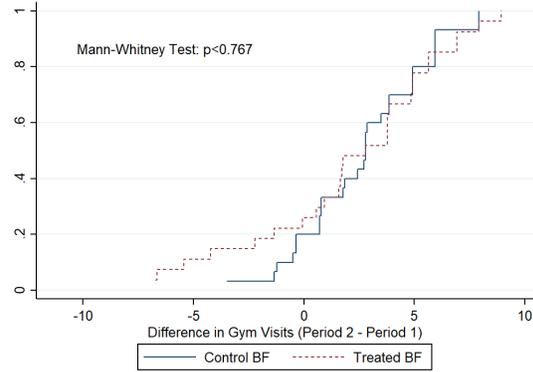


Table 1. Balance Across Participation and Treatment Status

	Balance on Participation		Balance on Treatment Status	
	Mean for Analysis Sample (1)	Difference for Dropouts (2)	Mean for Control Group (3)	Treatment Group Difference (4)
Age	18.212 [0.472]	0.047 (0.039)	18.179 (0.413)	0.058 (0.037)
Male	0.426 [0.495]	0.115** (0.039)	0.385 (0.488)	0.072* (0.040)
Black	0.023 [0.150]	0.012 (0.014)	0.034 (0.182)	-0.020 (0.013)
Hispanic	0.258 [0.438]	0.011 (0.034)	0.263 (0.441)	-0.009 (0.036)
Asian	0.194 [0.396]	0.061 (0.033)	0.191 (0.394)	0.006 (0.032)
Receives Financial Aid	0.546 [0.498]	0.001 (0.039)	0.573 (0.496)	-0.047 (0.041)
Self-Reported Exercise per Week	3.455 [2.012]	-0.244 (0.159)	3.429 (2.004)	0.033 (0.164)
Pre-Treatment Rec Center Visits per Month	2.228 [4.327]	-0.109 (0.335)	2.190 (4.230)	0.160 (0.352)
Treated Best Friend	0.578 [0.494]	--	0.584 (0.494)	-0.010 (0.040)
Sample size	612	839	262	612

Notes: In a series of balance regressions, we regress each of the observable student-level characteristics on an indicator equal to one if the student is in the analysis sample (Column 2). Dropouts participated in Wave I but not Wave II. We repeat these balance regressions but instead include an indicator equal to one if the subject is treated (Column 4). Standard deviations are in brackets. Robust standard errors in parentheses. ** p<0.05, * p<0.10

Table 2. Balance Across Participation Choice

	Mean Participation Rate for Those with No Treated Peer (1)	Difference in Participation for Those with a Treated Peer (TP) (2)
<u>Panel A: All Participants</u>		
TP = Any Treated Roommate	0.727 [0.446]	0.020 (0.031)
Sample Size	494	833
<u>Panel B: Participants who at least one person labels a Best Friend</u>		
TP = Any Treated Best Friend (friend-stated)	0.802 [0.399]	0.027 (0.038)
Sample Size	162	450

Notes: In two separate regressions, we regress the indicator for being in the analysis sample on an indicator equal to one if the subject has at least one roommate who is treated (Panel A) and then again on an indicator equal to one if the subject has at least one friend-stated best friend who is treated (Panel B). A balance test for self-stated best friends is not included because we only observe this reference group for subjects who return to Wave II. Standard deviations are in brackets. Robust standard errors are in parentheses. ** p<0.05, * p<0.10

Table 3. Treatment and Best Friend Spillover Effects for Gym Visits

	Visits (1)	8+ Visits (2)
Treated Subject	4.004** (0.422)	0.466** (0.049)
Treated Best Friend	-0.260 (0.341)	0.009 (0.031)
Treated Subject x Treated Best Friend	1.275** (0.569)	0.062 (0.064)
<u>Effect of a Treated Best Friend</u>		
For a Control Subject	-0.260 (0.341)	0.009 (0.031)
For a Treated Subject	1.015** (0.456)	0.071 (0.056)

Notes: Best friend is defined as self-stated first best friend. Sample size is 1,224. All models include individual and period fixed effects. Standard errors are in parentheses and are clustered at the individual level. ** p<0.05, * p<0.10

Table 4. Treatment and Spillover Effects for Gym Visits for Best Friend Network

	(1)	(2)
Treated Subject	3.755** (0.492)	3.755** (0.492)
Treated Best Friend (BF)	-0.248 (0.390)	-0.447 (0.596)
Treated BF x Treated Good Friend (GF)		0.249 (0.554)
Treated Subjects x Treated BF	1.524** (0.623)	1.874* (0.981)
Treated Subjects x Treated BF x Treated GF		-0.453 (0.949)
<u>Treatment Effects for Treated Subjects</u>		
Treated BF	1.276** (0.486)	
Treated BF Only relative to Treated GF Only		1.427* (0.779)
Treated BF and Treated GF relative to Treated GF Only		1.223** (0.510)
Treated BF and Treated GF relative to Treated BF Only		-0.204 (0.771)

Notes: Best friend is defined as self-stated first best friend and good friend refers to self-stated second and third best friends. Sample size is 1,114 (vs. 1,224) as we have dropped 55 subjects with neither a treated BF nor GF. The omitted group is control subjects with a treated good friend and an untreated best friend. All models include individual and period fixed effects. Standard errors are in parentheses and are clustered at the individual level. ** p<0.05, * p<0.10

Table 5. Fraction in Each Alternative Peer Category

	Alternate Peer (AP) is Roommate	AP is Friend- Reported Best Friend
Treated Best Friend and Untreated AP	0.33	0.27
Treated AP and Untreated Best Friend	0.09	0.08
Treated BF and Treated AP (same person)	0.13	0.24
Treated BF and Treated AP (different people)	0.12	0.07
Neither Treated Best Friend nor Treated AP	0.33	0.34
Sample Size	612	612

Table 6. Treatment and Spillover Effects for Gym Visits for Alternative Peer Groups

	Alternate Peer (AP) is Roommate (1)	AP is Friend- Reported Best Friend (2)
Treated Subject	4.086** (0.464)	4.080** (0.439)
Treated Best Friend (BF) Only	0.085 (0.388)	-0.104 (0.427)
Treated Alternate Peer (AP) Only	0.379 (0.648)	-0.407 (0.889)
Treated BF and Treated AP (same person)	-0.333 (0.608)	-0.665 (0.386)
Treated BF and Treated AP (different people)	-0.572 (0.571)	-0.006 (0.830)
Treated Subject x Treated BF Only	0.631 (0.687)	0.467 (0.719)
Treated Subject x Treated AP Only	-0.341 (1.096)	-0.254 (1.266)
Treated Subject x Treated BF and Treated AP (same person)	1.698* (0.916)	1.976** (0.706)
Treated Subject x Treated BF and Treated AP (different people)	2.218** (0.914)	1.453 (1.342)
<u>Effect of a Treated BF and/or Treated AP for a Treated Subject</u>		
Effect of Treated BF and Untreated AP	0.715 (0.568)	0.363 (0.578)
Effect of Treated AP and Untreated BF	0.038 (0.885)	-0.661 (0.902)
Effect of Treated BF and Treated AP (same person)	1.365** (0.685)	1.311** (0.591)
Effect of Treated BF and Treated AP (different people)	1.646** (0.714)	1.448 (1.054)

Notes: Best friend is defined as self-stated first best friend. Sample size is 1,224. The omitted group is a control subject with no treated alternative peers with an untreated self-stated best friend. All effects reported in the bottom panel are relative to a treated subject with an untreated AP and untreated self-reported BF. All models include individual and period fixed effects. Standard errors are in parentheses and are clustered at the individual level. ** p<0.05, * p<0.10

Table 7. Treatment and Spillover Effects for Gym Visits – OLS and IV Results

	First Stage: Best Friend Visits
<u>Panel A: First Stage</u>	
Treated Subject	-0.161 (0.314)
Treated Best Friend	3.657** (0.419)
Treated Subject x Treated Best Friend	0.871 (0.568)
<u>Effect of a Treated Best Friend</u>	
For a Control Subject	3.657** (0.419)
For a Treated Subject	4.528** (0.384)
F-Statistic for Weak Identification (K-P)	36.628
<u>Panel B: IV Estimates</u>	
	IV: Own Visits
Treated Subject	3.204** (0.771)
Best Friend Visits	-0.071 (0.096)
Treated Subject x Best Friend Visits	0.324** (0.154)
<u>Effect of Best Friend Visits</u>	
For a Control Subject	-0.071 (0.096)
For a Treated Subject	0.253** (0.113)

Notes: Best friend is defined as self-stated first best friend. The sample size is 1,224. Standard errors are in parentheses and are clustered at the individual level. Stock-Yogo (2005) critical value for 10% maximal IV size is 7.03. ** p<0.05, * p<0.10

Table 8. Number of Simultaneous Visits with Best Friend During Treatment Period for Various Time Windows

	1 minute (1)	5 minutes (2)	10 minutes (3)	15 minutes (4)	20 minutes (5)
Treated Subject	0.077 (0.115)	0.079 (0.122)	0.073 (0.153)	0.079 (0.160)	0.079 (0.163)
Treated Best Friend	0.141 (0.122)	0.129 (0.131)	0.116 (0.159)	0.108 (0.167)	0.114 (0.167)
Treated Subject x Treated Best Friend	0.887** (0.206)	0.964** (0.211)	0.981** (0.232)	1.020** (0.240)	1.049** (0.243)
<u>Effect of a Treated Best Friend</u>					
For a Control Subject	0.141 (0.122)	0.129 (0.131)	0.116 (0.159)	0.108 (0.167)	0.114 (0.167)
For a Treated Subject	1.028** (0.166)	1.092** (0.166)	1.097** (0.169)	1.128** (0.173)	1.163** (0.176)

Notes: Best friend is defined as self-stated first best friend. Sample size is 1,224. All models include individual and period fixed effects. Standard errors are in parentheses and are clustered at the individual level. ** p<0.05, * p<0.10

Table 9. Treatment and Spillover Effects for Gym Visits by Pre-Period Gym Visit Behavior

	Subject No Visits		Subject No Visits BF Some Visits		Subject Some Visits BF No Visits		Subject Some Visits BF Some Visits	
	8+ Visits BF No Visits	8+ Visits (2)	Visits (3)	8+ Visits (4)	Visits (5)	8+ Visits (6)	Visits (7)	8+ Visits (8)
Treated Subject	4.236** (0.545) [0.00]	0.401** (0.060) [0.00]	5.767** (1.003) [0.00]	0.711** (0.119) [0.00]	2.452* (1.054) [0.02]	0.349** (0.124) [0.00]	4.092** (0.965) [0.00]	0.600** (0.128) [0.00]
Treated Best Friend	0.063 (0.307) [1.00]	-0.018 (0.017) [0.95]	-0.503 (0.785) [1.00]	-0.032 (0.074) [0.95]	-1.399 (1.045) [0.72]	-0.097 (0.104) [0.95]	0.096 (1.037) [1.00]	0.200 (0.104) [0.22]
Treated Subject x Treated Best Friend	1.898** (0.655) [0.02]	0.234** (0.076) [0.01]	-0.038 (1.224) [1.00]	-0.246 (0.153) [0.32]	0.756 (1.573) [1.00]	0.033 (0.177) [0.85]	-0.721 (1.419) [1.00]	-0.230 (0.176) [0.38]
<u>Effect of a Treated Best Friend</u>								
For a Control Subject	0.063 (0.307) [1.00]	-0.018 (0.017) [0.95]	-0.503 (0.785) [1.00]	-0.032 (0.074) [0.95]	-1.399 (1.045) [0.72]	-0.097 (0.104) [0.95]	0.096 (1.037) [1.00]	0.200 (0.104) [0.22]
For a Treated Subject	1.962** (0.579) [0.00]	0.217** (0.074) [0.01]	-0.542 (0.939) [1.00]	-0.278 (0.134) [0.11]	-0.643 (1.176) [1.00]	-0.065 (0.143) [1.00]	-0.625 (0.969) [1.00]	-0.030 (0.141) [1.00]
Sample Size	616	616	188	188	210	210	210	210

Notes: Best friend is defined as self-stated first best friend. Standard errors are in parentheses and are clustered at the individual level. Holm p-values that adjust for multiple hypothesis testing across many subgroups are reported in brackets; p-values are adjusted for the four subgroups analyzed. ** p<0.05, * p<0.10, and corresponds to the robust standard errors rather than the Holm p-values.

FOR ONLINE PUBLICATION

Appendix 1

Survey and Consent Forms

1. Consent Form
2. Survey
3. Informational Form for Treated Subjects

PURPOSE:

You are being asked to participate in a research study. The purpose of the study is to discern the effects of exercise on the health outcomes of individuals and possible spillover effects on social groups to which an individual belongs.

PROCEDURES:

If you decide to participate, you will be asked to fill out a questionnaire about your personal background. Your consent today will mean that your name will be included on a list to be distributed to all participants of this study later in the quarter, for a follow-up friendship survey. And your consent today will allow us to access your attendance records at the UCSB Recreation Center between September 2010 and June 2011.

Your participation today will probably take about 5 minutes. There will be about 900 subjects in this study.

RISKS:

There are no significant risks for your participation today.

BENEFITS:

There will be no substantial benefit to you from completing these measurements and filling out the questionnaire, other than the money that you receive.

CONFIDENTIALITY:

Any records of your identity will be securely stored. Furthermore, any data stored on computers will only match to your identity by a numeric code. The information that matches your code to your identity will not be stored on a computer. When we no longer need to match your identity to your records, we will destroy any records that could match your data with your identity. Absolute confidentiality cannot be guaranteed, since research documents are not protected from subpoena.

COSTS/PAYMENT:

You will receive \$5 cash for your participation today, which will be paid when your participation has concluded. If you begin participation today and decide to continue in the research, you will receive \$8 to fill out a follow-up survey in a week or two. If you complete both surveys, we will invite some of you to participate in a follow up study on incentives to exercise. If you begin participation today and decide not to continue in the research, you will receive \$3. If you withdraw now, you will not be able to participate in the other portions of this study.

EMERGENCY CARE AND TREATMENT FOR INJURY:

If you are injured as a direct result of research procedures, you will receive reasonably necessary medical treatment at no cost. The University of California does not provide any other form of compensation for injury.

RIGHT TO REFUSE OR WITHDRAW:

You may refuse to participate and still receive any benefits you would receive if you were not in the study. You may change your mind about being in the study and quit after the study has started.

QUESTIONS:

If you have any questions about this research project or if you think you may have been injured as a result of your participation, please contact:

John Hartman (805-893-7309) or Philip Babcock (805-893-4823) or send an e-mail to: hartman@econ.ucsb.edu and/or babcock@econ.ucsb.edu.

If you have any questions regarding your rights and participation as a research subject, please contact the Human Subjects Committee at (805) 893-3807 or hsc@research.ucsb.edu. Or write to the University of California, Human Subjects Committee, Office of Research, Santa Barbara, CA 93106-2050

Signature

Name

Perm #

QUESTIONNAIRE

Please note that in order to participate in this study, you need to be a UCSB student living in Santa Catalina Residence Hall. I would like to ask you a few questions regarding some of your characteristics. Thank you.

NAME _____ Perm Number _____ Dormitory and room _____

Please list two e-mail addresses that we can contact you at:

Local phone number that we can contact you at: _____

1. What is your class standing Freshman Sophomore Other (specify) _____
2. What is your gender? Male _____ Female _____
3. How old are you? _____ years
4. What is your birth date? Month _____ Day _____ Year _____
5. Do you have at least one parent with a bachelor's degree?
Yes _____ No _____
6. Are you of Hispanic and/or Latino ethnicity? Yes No
7. What is your race?
_____ African-American
_____ Asian
_____ Caucasian
_____ Pacific Islander
_____ Other (Please specify _____)
8. How many times per week do you moderately or vigorously exercise for 30 minutes or more?
0 less than 1 1 2 3 4 5 6 7 more than 7
9. Did you qualify for financial aid [Any of the following: Federal Pell grant, Federal Supplemental Educational Opportunity Grants (SEOG), Academic Competitiveness Grant (ACG), Federal work study, Federal subsidized loan, Federal unsubsidized loan, Cal grant, University of California Grant]? Yes No

PURPOSE:

You are being asked to participate in a research study. The purpose of the study is to discern the effects of exercise on the health outcomes of individuals and possible spillover effects on social groups to which an individual belongs.

PROCEDURES:

If you decide to participate in this research follow-up to your recent participation in filling out a questionnaire, you will be given the opportunity to receive \$60 by exercising. In order to receive \$60, you will need to exercise at the UCSB Recreation Center (commonly known as the Rec Center) a minimum of eight different days over a four week period (**FEBRUARY 6 TO MARCH 6, 2011**). In order for any exercise session to count, you simply need to make sure that the computer at the front desk of the Recreation Center acknowledges your presence each day you exercise. You will receive the \$60 only if you make the eight visits to the UCSB Recreation Center as described above, but you do not have to make any visits if you do not want to.

RISKS:

Exercise has potential risks and benefits. Before starting any exercise program, you may want to consider contacting a doctor or other professional qualified to help determine what types of exercise are appropriate for you. When exercise is tailored to your physical condition and health, the gains from exercise usually outweigh the costs. Please also note that pregnancy may complicate the type and amount of exercise that you need. If you are pregnant or plan on becoming pregnant in the next two months you may not participate in this exercise study.

BENEFITS:

There will be no substantial benefit to you from completing these measurements and filling out the questionnaire, other than the money that you receive.

CONFIDENTIALITY:

Any records of your identity will be securely stored. Furthermore, any data stored on computers will only match to your identity by a numeric code. The information that matches your code to your identity will not be stored on a computer. When we no longer need to match your identity to your records, we will destroy any records that could match your data with your identity. Absolute confidentiality cannot be guaranteed, since research documents are not protected from subpoena.

COSTS/PAYMENT:

Again, you will receive \$8 for filling out a friendship survey. You will also be given the opportunity to receive an additional \$60 by exercising. In order to receive \$60, you will need to exercise at the UCSB Recreation Center (commonly known as the Rec Center) a minimum of eight different days over a four week period.

EMERGENCY CARE AND TREATMENT FOR INJURY:

If you are injured as a direct result of research procedures, you will receive reasonably necessary medical treatment at no cost. The University of California does not provide any other form of compensation for injury.

RIGHT TO REFUSE OR WITHDRAW:

You may refuse to participate and still receive any benefits you would receive if you were not in the study. You may change your mind about being in the study and quit after the study has started.

QUESTIONS:

If you have any questions about this research project or if you think you may have been injured as a result of your participation, please contact: John Hartman (805-893-7309) or Philip Babcock (805-893-4823) or send an e-mail to: hartman@econ.ucsb.edu and/or babcock@econ.ucsb.edu.

If you have any questions regarding your rights and participation as a research subject, please contact the Human Subjects Committee at (805) 893-3807 or hsc@research.ucsb.edu. Or write to the University of California, Human Subjects Committee, Office of Research, Santa Barbara, CA 93106-2050

Signature

Name

Perm #

Appendix 2

Table A1. Treatment and Spillover Effects for Gym Visits for the Friend Network

	(1)	(2)
Treated Subject	4.004** (0.422)	3.042** (1.237)
Treated Best Friend (BF)	-0.260 (0.341)	-1.165 (0.896)
Fraction of Friends Treated		-0.723 (0.935)
Treated Subject x Treated BF	1.275** (0.569)	1.343 (1.770)
Treated Subject x Fraction of Friends Treated		1.544 (1.988)
Treated BF x Fraction of Friends Treated		0.073 (2.823)
Treated Subject x Treated BF x Fraction of Friends Treated		0.773 (0.477)
<u>Effects of Treated BF and Marginal Effects of Fraction Friends Treated</u>		
Control Subject with Treated Best Friend	-0.260 (0.341)	
Control Subject, Treated BF: Effect of Fract. Treated		-0.650 (3.254)
Control Subject, Control BF: Effect of Fract. Treated		-0.723 (0.935)
Treated Subject with Treated Best Friend	1.015** (0.456)	
Treated Subject, Treated BF: Effect of Fract. Treated		2.479 (1.643)
Treated Subject, Control BF: Effect of Fract. Treated		0.862 (1.755)

Notes: Best friend is defined as self-stated first best friend. Sample size is 1,224. All models include individual and period fixed effects. In column 2, fraction of friends treated excludes the best friend. Standard errors are in parentheses and are clustered at the individual level. ** p<0.05, * p<0.10

Figure A1. Weekend Rec Center Visits (per/month)

