Chapter 6

Experimental Approaches to the Diffusion of Norms

David W. Nickerson

Many theories of social capital contend that dense ties within networks develop norms that govern the behavior of group members also establishing surveillance and sanctioning systems for violations of those norms. Compliance with those norms provides the basis for trust in the community as behavior become more predictable. Coleman (1988), Burt (2005) and many other scholars argue that networks with closure are sources of social capital precisely because densely connected groups transmit and reinforce group norms. Thus, social networks are perceived as critical in the diffusion and reinforcement of important civic and other norms. Unfortunately, selection into social networks makes it extremely difficult to isolate the unique role played by the network in the formation of norms. Specifically, how does one disentangle the influence of the network over time from that of the traits that influenced the person to join the network in the first place?

Randomized field experiments offer a method for measuring the role of these two processes. Researchers can pursue three different experimental strategies to measure norm creation in the field: *create* a randomized social network; randomly manipulate *interactions within* an existing network; or, trace the effect of an *exogenous shock* through the network. This chapter briefly explains the logic motivating experiments and provides an example of each type of experiment to study norm formation within social networks.[1]

*****

Friends, family members, co-workers and neighbors influence one another's actions and attitudes. Neighbors tend to vote in blocs[2] and children often follow their parents' lead in political behaviors and beliefs.[3] These phenomena lead some to believe that interpersonal influence is obvious and ubiquitous. This very ubiquity, however, makes accurate detection of interpersonal influence a complicated endeavor.

For instance, neighborhoods often exhibit a remarkable homogeneity with regards to aesthetics, social status, and even norms. Social pressures may lead to conformity regarding lawn care standards, but it is also possible that the family chose to move into the neighborhood because they felt comfortable with the existing norms. Casual discussions about politics can create agreement on political matters, but people living in close proximity often possess similar incomes and, thus, share similar material interests that shape political views. Nearly every example of norm formation within social networks possesses an equally plausible counter explanation involving selection. Randomized experiments offer a method of isolating the effects of interpersonal influence.

[1] Each experiment is presented at length in Nickerson 2005.
[2] V.O. Key offers one of the earliest and most famous examples of context effects in voting.
[3] Newcomb and Svehla established this using modern survey techniques in 1937 (Newcomb and Svehla 1937), but the observation dates back at least as far as Plato. Jennings and Niemi (1974) provide a thorough analysis of this topic.

The Benefits of Randomization

For our purposes, an experiment is any random application of a factor on a pool of subjects. Classic experiments randomly divide subjects into a treatment group, which receives an intervention, and a control group, which does not.[4] Because the division is random, the treatment group and control group should be comparable. In other words, the likelihood that a subject receives the intervention is not affected by their age, income, education, gregariousness, curiosity, or any other trait. An experiment has been correctly designed when a subject's assignment to treatment or control provides no information about their demographic or other characteristics.

The chief attraction of experiments is the ease with which they can isolate unique influences of the manipulated treatment. The empirical leverage afforded by experiments can be illustrated by thinking of one person's behavior as a function of another person's behavior and outside factors. Equation 1 models the effect of person 1's behavior on person 2 where A_2 represents the attitudes or actions of person 2, c is a constant, A_1 the attitudes or actions of person 1, X are the variables that can be measured and controlled for, $g(H)$ are the known causes of both people's attitudes, and ε is the idiosyncratic causes of person 2's behavior:

$$A_2 = c + \alpha A_1 + \beta X + g(H) + \varepsilon \quad (1).$$

[4] Application of the treatment can be split as many ways as the researcher desires. Medical researchers often experiment with different doses of the same drug. Lotteries like Powerball could be thought of income experiments with the range of winnings constituting the different treatment groups (and the losers serving as the control group). Nor does the treatment need to involve a single factor. Factorial design can easily accommodate interactions between multiple treatments. For instance, suppose one wanted to know whether citizens are mobilized by leaflets, a knock on the door, or a combination of both. To test this, registered voters in a neighborhood could be randomly assigned to receive only a leaflet, only a knock on the door, both the leaflet and the knock on the door, or no contact at all from the campaign. The voter turnout rates between the four groups could then be compared to determine which technology is more cost effective increasing turnout. This type of protocol could be extended to any number of treatments. Sample size is the only limitation upon the use of the technique.

Because H is unmeasured and correlates with A_1, $E[\varepsilon \mid A_1] \neq 0$; as a result, estimates of α may be biased using most observational techniques.

But suppose that an experimenter devised a study where subjects were randomly divided into a treatment group that was exposed to the opinions of their neighbors ($A_1 = 1$), and a control group that was never informed of how neighbors felt ($A_1 = 0$). We could then apply Equation 1 to both the treatment group, denoted by subscript T, and the control group, denoted by subscript C. Equation 2 presents the difference between the treatment and control groups:

$$\begin{aligned} A_{2C} &= c_C + \alpha A_{1C} + \beta X_C + g(G_C) + \varepsilon_C \\ \underline{A_{2T}} &\underline{= c_T + \alpha A_{1T} + \beta X_T + g(G_T)} + \varepsilon_T \\ A_{2T} - A_{2C} &= c_T - c_C + \alpha(A_{1T} - A_{1C}) + \beta(X_T - X_C) + g(G_T) - g(G_C) + \varepsilon_T - \varepsilon_C \end{aligned} \qquad (2).$$

Because the treatment and control groups are randomly assigned, the manipulated factor A_1 should be the only independent variable that differs between the two groups. With sufficiently large sample sizes the observed, unobserved, and contingent factors balance out. This is expressed mathematically as $E(c_T - c_C) = 0$, $E(X_T - X_C) = 0$, $E[g(G_T) - g(G_C)] = 0$, and $E(\varepsilon_T - \varepsilon_C) = 0$. Thus, the treatment is dichotomous, equation 2 can be reduced to:

$$A_{2T} - A_{2C} = \alpha(A_{1T} - A_{1C}) = \alpha \qquad (3)$$

In words, to estimate the effect of one person's opinions on another person's opinions, α, the experimenter simply needs to subtract the mean of the control group from the mean of the treatment group when a randomized experiment is conducted.[5] The experiment's design and

[5] Experimental analysis is not devoid of assumptions. However, the analysis above reduces the number of assumptions required in the analysis. No assumptions were required about the process that generated the data, the distributions of the variables, or the relationship between the variables. Other assumptions were still necessary. For instance, equation 2 could not be simplified without $E(c_T - c_C) = 0$ and $E(\varepsilon_T - \varepsilon_C) = 0$. That is, it is presupposed that the treatment and control groups have the same intercepts and residual errors on average. However, these two assumptions rest upon the Law of Large Numbers and executing the protocol successfully. It should also be pointed out that observational studies make the same assumptions, so experimental interventions are

careful implementation, thereby allow the simple subtraction of two averages to yield an unbiased estimate of a particular form of interpersonal influence.[6]

Laboratory versus Field Experiments

Laboratory experiments are commonly employed tool in social science. They have examined the effect of discussing news with others (e.g., Iyengar and Kinder 1987), the role of personal information in trust games (e.g., Wilson and Musick 1997), exposure to discordant ideas on political tolerance (e.g., Mutz 2001), and the spread of rumors (e.g., Allport and Postman 1940).

Unfortunately, such studies remove subjects from the social networks the experimenter is trying to assess. Because the setting is artificial, findings may or may not translate to the original networks. While such studies provide interesting insights into interactions in anonymous settings (of which there are many), they appear to be severely limited in their ability to deliver findings about the original social networks. Ultimately, the extent to which laboratory findings translate to actual neighborhoods needs to be tested. If the goal is to describe interactions in settings in which subjects exist on a day-to-day basis, field experiments are essential to confirm and extend results from the lab.

There is no obvious method of making organic social networks amenable to experimental manipulation. In many instances, the task will prove impossible. However, in principle,

---

not adding to the presuppositions of the model. The uncertainty associated with the experimental estimate of α is due to sampling error, not model misspecification.

[6] The simplicity of the unbiased estimator of α in equation 3 does not imply that more complicated econometric estimators cannot be used. In fact, modeling assumptions are better satisfied by experimental rather than observational data. Where observational data requires covariates, multiple stages, and modeling assumptions to claim $E[\varepsilon \mid A_1] = 0$, the condition holds by construction of the variable A_1 in experiments. Adding covariates to the analysis reduces variance and thereby shrinks standard errors, but covariates are unnecessary to achieve unbiased estimates of α because omitted variable bias is not a concern.

randomized field experiments can study interpersonal and social influence in three ways: 1) *create* a randomized social network; 2) control the *flow of communication* within an existing social network; or 3) introduce *exogenous shocks* into the network and trace the effect of the shock.[7] Examples of each one of these tactics are presented in order.

Creating and Measuring a Randomized Network

Observational studies of neighborhoods such as Oliver (2001) and Huckfeldt and Sprague (1995) may report biased estimates of influence and participation because subjects self-selected into the communities and neighborhoods where they live. But what if it were possible to randomly assign residents of South Bend to neighborhoods? Random assignment would solve the confounding influence of selection, and the models Huckfeldt and Sprague analyze would yield unbiased estimates of the influence of political context. Obviously, such an experiment would be unfeasible and unethical.

However, randomized assignment may be possible *within* specialized communities. There are many settings in which persons are assigned to groups: recruits to camps and barracks in the military; poor families to housing developments, apartment buildings, and specific apartments (see Katz, Kling and Liebman 2001); prisoners to prisons, cells and blocks; and college students to dorms and rooms (see Sacerdote 2001; Winston and Zimmerman 2003; Stinebrickner and Stinebrickner 2001; Hoel et al. 2004; McEwan and Soderberg 2004; Marmaros and Sacerdote 2002; Duncan et al. 2003). The cases of college students and prisoners are particularly interesting because the governing authority not only assigns rooms but also room- or

[7] These three protocols are not mutually exclusive; in principle, a researcher could design an experiment where communication about a controlled external event takes place within a randomly constructed social network. However, given the difficulty in executing each experiment on its own, combining the different types may be overkill.

cellmates. After assignments are made, people explore the social topography and form relationships with those around them.[8]

Usually assignments within these communities use a conscious decision-making process. Many colleges ask incoming students to answer surveys of varying lengths to determine the compatibility of roommates and create demographic balance within dorms. But in the instances where assignments are made randomly (not just haphazardly), an experiment has been conducted and all a researcher has to do is collect the data. Because all the persons of a certain type have been randomly distributed, there is no risk that civically engaged persons opted into the same apartment building or that all violent criminals are in a particular cellblock. If interpersonal influence effects are strong enough, pockets of norms, behaviors, and opinions will form beyond what is attributable to random chance. If the context effects are weak, then any pockets of homogeneity will probably be the result of sampling variance (i.e., a group of Republicans placed together by random chance) and estimates of influence will not reject the null hypothesis of individuals being unaffected by their roommates.[9]

It is important to note what is measured when assignments are controlled within a social network. Such experiments measure context effects and not individual interactions. While neighbors (and perhaps roommates) are externally imposed, the subject remains free to select the people with whom they interact the most – their friends. A regression analysis considering the influence of the norms in an apartment building will yield accurate estimates of the effect of these norms on the behavior and beliefs of the individual subject. In contrast, a survey comparing the attitudes of friends will be biased because of self-selection and unobserved

[8] This set up differs from lab experiments in that the subjects are living within the social context as opposed to operating in it for brief periods of time. The interactions are not artificial by definition because the constructed framework really is the social network in which the subject resides.
[9] To be more precise, the null hypothesis will not be rejected 95% of the time.

heterogeneity. The randomization of one facet of social interaction does not purge the entire process of bias.[10]

An experiment that used randomized social networks to study the spread of norms was conducted during the 2002-2003 school year at a college that randomly assigned new students to housing (see Nickerson 2005, chapter 3). The school is typical of residential campuses across the country. It is a somewhat selective, mid-sized private school in a suburban East Coast setting. Unlike students at Dartmouth, Reed and Williams College (i.e., the sites of the grade point average studies), the students at this college were not the academic elite. Racially, the student body mirrors broader society, however, with tuition, room and board costing more than $30,000, the students come from wealthier than average families.

Students were randomly assigned to rooms within suites of dormitories. Rooms generally contained two students and suites consisted of four rooms sharing a living space and bathroom. Given the relatively intimate accommodations, one could reasonably expect behavioral contagion among both roommates and suitemates. At the beginning of the school year, there should be no correlation between the behaviors and beliefs of randomly selected roommate pairs, as every student has an equal chance of being placed with a given roommate of the same gender.[11]

The predisposition of roommates (and suitemates) constitutes the treatment to which a student was exposed. Some students had apathetic roommates with no real political opinions,

---

[10] It is perfectly acceptable for the randomization to occur within strata. There may be important policy concerns that trump the researcher's desire for a purely random sample: the severity of the crime and the predicted risk of future violence help to determine the assignment of prisoners to wards; rooms (and sometimes floors) of dorms are gender segregated; and, the military provides different training and assignment to specialized troops such as intelligence or communications. While such stratification needs to be accounted for and limits the conclusions to be drawn, it in no way invalidates the experiments as a whole. The researcher simply needs to examine the context effect within the strata created rather than in the system as a whole. The net effect is conducting several concurrent experiments on slightly different populations.

[11] It is possible that men and women differ with regards to a behavior or a belief on average. As a result, it is important to control for gender when studying roommate pairs.

while others will be presented with strident Democrats or Republicans. Nonrandom clusters of opinions and behaviors at the end of the school year (i.e., roommates possess correlated behaviors and beliefs) would provide evidence of influence.

To measure the influence, one can model student *i*'s end of year attitudes (A_{2i}), which are a function of her attitudes in the beginning of the year (A_{1i}) and the starting opinions of her roommates (A_{1R_i}) and suitemates (A_{1S_i}). To account for the single sex nature of rooms and suites, a dummy variable for gender is also included. Equation 5 summarizes the regression equation to be used where *c* is the constant and ε_i represents the unobserved, idiosyncratic causes of a student's opinion at the end of the year.[12]

$$A_{2i} = c + \beta_1 A_{1i} + \alpha_1 A_{1R_i} + \alpha_2 A_{1S_i} + \beta_2 Female + \varepsilon_i \qquad (5)$$

Data for the study came from a 28 question, two-wave panel survey administered during the second day of freshman orientation and the week prior to final exams in the spring. Roughly half of the entering class took both the first and second wave of the survey. However, 45% of the students did not complete the full year of study and not all respondents had roommates and suitemates complete the survey. As a result, the sample of students answering both surveys with pre-survey information for roommates and suitemates is relatively small (N=230). Fortunately, beginning-of-the-year attitudes explain most of the variance in end-of- year-attitudes, so the study has sufficient statistical power to detect influence.[13] Questions probed a range of political

[12] Following the lead of Sacerdote (2001), when a student has two roommates, the views of both roommates will be averaged together to calculate A_{1R_i}. Similarly, A_{1S_i} represents the average views of person *i*'s suitemates. Averaging loses information but in no way biased the analysis.

[13] No meaningful correlation was found between pre-existing attitudes among either roommates or suitemates.

beliefs ranging from broad ideology to sentiments regarding specific groups or attitudes about particular programs. Table 1 presents the results for those questions exhibiting influence.[14]

Insert Table 1 about here

The bottom line of Table 1 is that roommates and suitemates can shift a person's opinions, but the shift is neither radical and nor does it necessarily breed like-mindedness. Imagine two students who begin the year indifferent to homosexuals (i.e., rating of 0). One student is placed with a roommate who has an unfavorable view of homosexuals (i.e., -2), the other student with a roommate with a favorable view of homosexuals (i.e., +2). By the end of the year, our results predict that one student will have a slightly unfavorable view of homosexuals (i.e., -0.5), the other a slightly favorable view (i.e., +0.5).

Even more striking is the fact that a student's attitude at the beginning of the year is by far the best predictor of the student's opinion at the end of the school year.[15] The coefficients are much larger than the peer effect coefficients, with most in the range of 0.35 to 0.56 and statistically significant. For the homosexual question, the pre-test coefficient is over *twice* as high as the roommate coefficient and the coefficient for the pot smokers question was over *three* times larger than the suitemate one.

Some behaviors do seem to shift somewhat based on roommate and suitemate influence. Few would be surprised that students whose roommates enjoy discussing politics also report discussing politics (though they are no more interested in politics at the end of the year). Similarly, students whose suitemates regularly read the newspaper or watch television are more

[14] The Q-statistic testing the homogeneity of treatment effect across questions is firmly rejected.
[15] This finding also validates the survey instrument; suggesting that the responses are measures of genuine opinion and not an artifact of taking the survey itself.

likely to consume such readily available media. However, even opinion questions on the legal drinking age and marijuana may capture behaviors. Roommates who procure illicit substances may alter attitudes simply by lowering barriers to usage.

Attitudes concerning homosexuality are the only exception to this pattern. A possible explanation is that, freed from direct parental oversight, college represents the first opportunity for young people to encounter homosexuality. The intimate living conditions with non-family members may increase the relevance of the topic. The combination unfamiliarity and importance probably creates an environment where interpersonal influence is strong.

The results of the study are not surprising, but they are instructive. By far the best predictors of a student's attitudes at the end of the year are her attitudes at the beginning of the year. Moreover, no peer effects whatsoever were detected on bedrock political attitudes such as partisan identification, taxes, and welfare. The views of roommates and suitemates matter, but only on behavioral, unfamiliar or highly salient topics and the effects are seen only on the margins.

Randomized Structured Interactions within an Existing Network

If answers about common residential life are desired, one must study behavior and interactions within the existing social network. The second experimental design seeks to do precisely this by controlling and measuring interaction between subjects in a current network. The literature on interpersonal influence within a network posits two kinds of mechanisms for convergence in attitudes and actions. In the first, friends and neighbors explicitly transfer information and ideas. This fits the classic Norman Rockwell image of neighbors talking over the fence or meeting at the corner store. Politics, for instance, may be touched upon only briefly, but communication of ideas around politics is explicit.

The second mechanism involves a diffusion of norms and expectations among peers. Friends and neighbors do not talk about politics as such, but subtle cues of approbation and disapproval are sensed by individuals. Eventually, residents internalize community norms without resorting to formal processes. The two mechanisms, both of which involve learning/conditioning, are mutually reinforcing. In Chapter 5 of this book, David Campbell (2007) uses observation data to examine the effect of both mechanisms in creating and reinforcing norms on voter turnout.

When considering these two approaches, the first is easier to tackle experimentally because its communication is direct and explicit.[16] The researcher simply needs to force members of the network to comply with a protocol that dictates the content of conversations. This can be done by recruiting members to participate in the experiment. The model is similar to how Huckfeldt and Sprague (1995) ask subjects to provide the names of people with whom they discuss politics. Huckfeldt and Sprague (1995) asked subjects questions about *past* interactions with the friends listed (and then interviewed the friends in the same fashion).

In the type of experiment described here, the volunteers are asked to *initiate* communication with a select set of friends. The list of friends is randomly divided into the various treatment groups and the volunteer is then instructed to carry out the protocol. The type of treatment can vary by topic (e.g., donations to candidates, voting, volunteerism, trimming the lawn, or simply holding an opinion) or by mode of delivery (e.g., face-to-face conversations, phone calls, or e-mail). The important thing is that the researcher – not the volunteer –

[16] Experiments testing the gradual diffusion of norms are better off examining constructed social networks.

determines who is contacted and how. This ensures the treatment and control groups do not differ in a subtle way.[17]

An experiment controlling the flow of information through social networks was conducted during the 2002 Congressional election. Its goal was to determine the extent to which explicit appeals from friends and coworkers could increase voter turnout. Past field experiments had demonstrated that a brief face-to-face conversation with a stranger could increase vote turnout by 8 to 10 percentage points (Gerber and Green 2000; Green, Gerber, and Nickerson 2003). Given reputation effects and increases in social pressure, the hypothesis was that encouragement to vote from friends should be even more effective. Laboratory experiments have found this to be the case across a wide range of settings (e.g., Davis and Rusbult 2001; Mackie, Worth, and Asuncion 1990; Brock 1965) and political organizations have acted accordingly (e.g., the AFL-CIO's labor-neighbor program; MoveOn.org's friendship networks; the GOP peer-to-peer mobilization in 2004). There is good reason to believe that organizations can harness social networks to change norms of voter turnout.

An alternative hypothesis is that friends who are willing to volunteer time to encourage voting may operate in social circles where few people can be mobilized. The majority of the volunteer's friends will already vote, so it will be impossible to motivate them to do more. Another potential issue is that serial abstainers residing in social networks where voting is the

[17] One problem that may arise is treatment contagion where members of the control group interact with members of the treatment group and are exposed to some portion of the treatment. The problem of contagion can be bypassed in two ways. First, the researcher can simply select a set of subjects who are unlikely to infect one another. The distance need not be spatial; it could be along any dimension, such as social or temporal. The second solution is to randomly vary the density of the treatment (i.e., the subjects in one set receive a high dosage of the treatment while subjects in another set receive a low dosage). If the treatment is contagious, then the estimated treatment effects will be smaller for populations with a high percentage of subjects assigned to the treatment group than populations with a low percentage of subjects provided the treatment. While this will not eliminate contamination of the control group by the treatment group, it does provide a means of detecting the source of error.

norm probably have a reason for not voting and will resist strongly. They may be ineligible or simply detest politics.

An experiment was designed to test these hypotheses; volunteers were asked to list friends, neighbors, and coworkers they thought were unlikely to vote in the Congressional election. These individuals were then randomly assigned to be either contacted by the volunteer about voting or left alone.[18] After the election, the volunteers reported the names of the people they spoke with and official voter files were checked to determine who had voted. Analyses were run to determine whether subjects assigned to the treatment or control group were more likely to vote.

Volunteers were recruited from two sources. The first source was a 501(c)3 organization named Vote for America that mobilized voters via peer-to-peer networks. The organization had been active during the 2000 Presidential election in Rhode Island. It was eager to evaluate its model in an area where its volunteers were experienced, well-trained, and received robust organizational oversight and support. Thirty-one volunteers were recruited, who listed 481 friends to be mobilized. One-fourth of each volunteer's list was assigned to the control group and the remaining three-fourths were targeted for mobilization.

The second sample of volunteers was recruited on-line through email appeals. Ultimately, 65 volunteers successfully completed the protocol and lived in regions where voter databases were readily accessible. These 65 volunteers listed a total of 374 friends and neighbors; these were evenly divided into treatment and control conditions. Given the lack of training and support for Internet volunteers, one might expect the volunteers recruited over the

[18] Randomization tests indicated that the treatment and control were extremely similar with regards to observable characteristics. The only noticeable difference is that the control group for the internet volunteer sample was slightly less likely to be registered than the treatment group. However, this difference did not cross traditional thresholds of statistical significance and there is no reason to assume the randomization failed in any sense.

Internet to be less effective than Vote for America participants, so the two experiments are analyzed separately. Table 2 reports the results for both experiments.

Insert Table 2 about here

Neither experiment demonstrated any increase in voter turnout after accounting for slight imbalances between treatment and control groups. In fact, individuals who were approached by a friend regarding the election were an average 4.6 percentage points *less* likely to vote in the internet experiment and 3.2 percentage points *less* likely to vote in the Vote for America experiment. While a demobilization effect from peer-to-peer contact is extremely hard to believe, the experimental results strongly suggest that peer-to-peer mobilization is less effective than traditional mobilization from a stranger.

Given the consensus in the literature on the strength of personal ties, this null finding is surprising. Two explanations immediately suggest themselves. The first is that selection into homogenous social networks outweighs the importance of reputation and persuasion. Many people declined to participate in the experiment because "everyone I know votes." For such individuals, contacting members of their social circles will do little to increase turnout because everyone is planning to vote already. Participating volunteers were instructed to target people whom they thought might not vote. It might be possible that individuals who operate in social circles where voting is the norm but who regularly abstain cannot be mobilized.

A second explanation is that friends and neighbors actually are less effective at mobilizing voters. Perhaps developed trust between two individuals has narrow bounds and requests outside of these boundaries have little or no influence. The knock on a door from a friend to ask for civic participation may have seemed artificial. Why bring up voting after many years? The appeal may have even been perceived as insulting. If this explanation is true, then it is possible that in this case norms are transferred not through direct conversations but through subtext and social expectations. While the study is obviously small, it is suggestive of how interpersonal influence may need to be indirect in some situations.

Exogenous Shocks to a Network

It is not always possible to manipulate structured interactions within a network. An alternative strategy allows the researcher to focus on the influence of an external shock as it diffuses through the network. The exogenous shock design is versatile and can address a range of topics. The diffusion of information, the contagion of behaviors, and the passage of opinions through influence are all designed and analyzed the same way. It can also deal with almost any topic and whether the subjects should be informed, persuaded, indoctrinated, or something else.

However, attention to detail is paramount to avoid selection bias. The key step is to establish baseline rates for the behavior or belief in question. In order to measure change, the researcher needs to know what the subject's starting point is and then what the difference is.[19]

There are two primary methods for doing this. The more traditional one involves taking a measure of the network prior to the introduction of the exogenous shock (e.g., a survey). The results obtained after the intervention are then compared to this baseline. While surveys are not randomized experiments, they can offer a scientific mode of analysis in this setting. This is because the researcher is actively providing a treatment rather than passively measuring the effect of past interactions. Other researchers can replicate the pattern of diffusion measured by

[19] This point may sound pedantic and simplistic, but it is often ignored.

the researcher at different times and places, on different topics. This model of analysis is akin to a physicist dropping a rock in a pool and attempting to model the waves.[20]

The second model uses a placebo as the reference group. Two parallel exogenous interventions are made and they each serve as the other's point of reference. Both treatment and placebo need to meet two stringent requirements. First, the two interventions must not lead to the same outcome. If the treatment of interest is expected to improve political knowledge, then the placebo should be unrelated – even indirectly – to politics. In addition, the two treatments must travel through the network at roughly the same rate. A placebo that does not travel far from the point of its introduction would be a poor comparison for a treatment that spreads through a network quickly.

Selecting two treatments that meet both requirements can be difficult. The extent to which a design meets these two criteria is an empirical matter and can be difficult to gauge prior to conducting the study.[21] However, thinking through the realities can often give experimenters a sense of major differences. For instance, encouraging families to recycle paper is not a very good placebo for a treatment encouraging families to recycle aluminum because the two behaviors are likely linked. A rumor concerning police brutality is likely to spread through a neighborhood much more quickly than a piece of information on the mayor's electoral platform.

---

[20] A fabulous example of this protocol was provided by Dodd (1952). One in five people in a town of 950 people were told a six word slogan for a coffee company and were told that anyone able to repeat the slogan in two days would receive a free pound of coffee. The next day a plane dropped thirty thousand leaflets over the town. The leaflets stated that one in five people in town knew a six word slogan and anyone able to repeat the slogan would receive a free pound of coffee. The day after the massive literature drop, researchers knocked on every door in the town to see who knew the slogan (and hand out the coffee). Dodd conducted similar experiments in other cities and using different modes of mass communication.

[21] There is no reason why the two protocol designs cannot be combined. The panel framework has the advantage of mapping the network prior to the intervention so the system is more contained and easier to trace. The placebo design has the advantage of controlling for all temporal effects since both interventions take place at the same time. The strengths of each design cover the weaknesses of the other, so the dual design is optimal. The added complexity makes the execution of the protocol more difficult, but that is the only downside.

An experiment using the exogenous shock design was conducted during the 2002 Congressional Primaries in Denver and Minneapolis. The social network considered was very small – households with two registered voters – and the behavior traced was voter turnout. Various two-voter households in Denver and Minneapolis were randomly selected to receive face-to-face encouragement on either of two subjects: recycling or voting.[22] The mobilization effect of face-to-face canvassing is well established (see Gerber and Green 2000 or Gerber, Green, and Nickerson 2003), so it makes an excellent exogenous shock. The hypothesis is that people who answer the door and receive the Get Out The Vote (GOTV) message will turn out at higher rates than people who answer the door and receive the recycling message. But what happens to the person in the household who did not answer the door? How much of the boost in turnout from canvassing is passed on to that other person?

The design is notable in a few respects. First, nothing in the collection of data relies upon subject responses. The paid canvasser recorded who answered the door (if anyone) and the county clerk recorded who voted. Any measurement error is therefore rare random errors in transcription. Second, the placebo-controlled design offers a perfectly comparable reference group. Persons who answered the door are compared only with other persons who answered the door. Similarly, influence is examined only among persons whose cohabitants answered the door. The only difference between the groups is the treatment provided.[23] Finally, two voter households are convenient to work with for the experiment, but this analysis would be impossible with non-experimental data. Researchers often ignore selection with regards to city or neighborhood, but there is no avoiding the selection of spouses and roommates. The selection

[22] Households were also assigned to a control group that received no visit from the campaign. The purpose of the control group was simply to verify that the mobilization effect detected was genuine.
[23] The similarity of the treatment and placebo groups on observed characteristics is striking. Every randomization check supports this conclusion.

of cohabitants is not an issue in the slightest since the randomization of the treatment occurs only within the set of people who have selected each other. The experiment considers a unique and small social network, but it measures voter contagion very precisely.

When pooled together, the two experiments showed a 9.8 percentage point boost in turnout from being exposed to the GOTV message. In Denver, the campaign contacted 33% of the households in both the GOTV and recycling conditions for a total of 563 households receiving either the treatment or the placebo. Among the people who answered the door, those receiving the voting appeal were 8.6 percentage points more likely to vote than people answering the door and hearing the recycling message. In Minneapolis, 45% (394) of the target households were successfully contacted. The direct treatment effect was a robust 10.9 percentage points.

Surprisingly, over 60% of this boost in turnout was passed onto the other person in the household. Persons not answering the door in Denver households exposed to the voting appeal were 5.5 percentage points more likely to vote and in Minneapolis the boost in turnout was 6.4 percentage points over households assigned to the recycling message. Pooled together, we find registered voters in the household were 6 percentage points more likely to vote even though they were not directly exposed to the voting message. Table 3 uses two-stage least squares to pool the results and control for past voter history and demographic information.[24]

Insert Table 3 about here

[24] A bivariate probit model accounting for the dichotomous nature of voter turnout yields substantively identical results.

The magnitude of the effect is substantial and indicates that voting is a highly contagious behavior. A person who may have been 20% likely to vote in the primary becomes 80% likely to turn out should their housemate vote. A model where the spouse's or housemate's turnout is the only variable explains over 45% of the variance in a person's decision to vote.[25] These experiments suggest that the behavior of spouses and housemates has an extremely large impact on the decision to vote.

Conclusion

Similarities between pairs of people can be the result of a number of factors: interpersonal influence, similar pre-existing dispositions, structural incentives, a selection process that weeds out dissimilar individuals, exposure to identical exogenous factors, or social linkages. Unfortunately, "The observational data that would be necessary to disentangle each of these component parts fully and directly are unavailable, and very nearly unimaginable" (Huckfeldt and Sprague 1995, 164). Because of this, scholars who use observational data to try and estimate the effect of interpersonal influence, α , are often required to make unverifiable assumptions. In contrast, randomized field experiments require minimal assumptions about the process that generated the data or the relationship between the variables.

Part of the reason that experiments have not been used more in social science is that the questions asked are necessarily focused and therefore "small." The required control over the factor of interest limits the number of questions a researcher can answer in a given study and increases the difficulty in conducting research. This means studies on social capital, networks, and influence will need to be tightly focused, often dealing with one of the narrower slices or definition of social capital.

[25] The predictive accuracy is all the more astounding when the dichotomous nature of voting is taken into account.

The virtue of this approach is that experiments allow small questions to be answered accurately. The researcher is forced to use the incremental notion of intellectual progress found in the natural sciences. The three experiments described in this paper illustrate how one might apply experiments to social networks. The findings are interesting and informative, but readers anticipating the final word on interpersonal or influence will be sorely disappointed.

That said, the three experiments contained in this chapter paint a picture of a subtle and slow process of norm diffusion. While voter turnout is contagious within a household, explicit appeals from members of a social network to vote are wholly ineffective. For some reason, the effort to harness the power of the social network to encourage the norm of voting backfired. Similar limits were observed with regards to opinion contagion among first year college students. Contagion was observed only on a select few topics and behaviors, while individuals remained unmoved on the vast majority of questions. Moreover, the attitudinal contagion was highly constrained by pre-existing beliefs.

These results run counter to many observational studies where readers are lead to believe we live in a world where citizens are constantly buffeted by the attitudes and actions of neighbors. Every person is a party switcher who just happened to move into a particular neighborhood in a particular suburb. The truth is that the diffusion and formation of community norms is a slow process. A detailed and incremental experimental approach can unravel the bigger questions of just how social capital is formed under conditions of social closure in a community. It is my hope that the modest examples included in this chapter will spur other scholars to pursue more ambitious experiments in the field.

References

Allport, Gordon W. and Leo J. Postman. 1940. *The Psychology of Rumor*. New York: Henry Holt.

Brock, Timothy C. 1965. “Communicator-Recipient Similarity and Decision Change.” *Journal of Personality and Experimental Social Psychology* 1:650–54.

Davis, Jody L. and Caryl E. Rusbult. 2001. “Attitude Alignment in Close Relationships.” *Journal of Personality and Social Psychology* 81:65-84.

Dodd, Stuart C. 1952. “Testing Message Diffusion from Person to Person.” *Public Opinion Quarterly* 16:247-62.

Duncan, Greg J. Johanne Boisjoly, Dan M. Levy, Michael Kremer, and Jacque Eccles. 2003. “Empathy or Antipathy? The Consequences of Racially and Socially Diverse Peers on Attitudes and Behaviors.” Working paper, Harvard University.

Gerber, Alan S. and Donald P. Green. 2000. “The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment.” *American Political Science Review* 94:653-63.

Green, Donald P., Alan S. Gerber, and David W. Nickerson 2003. “Getting Out the Vote in Local Elections: Results from Six Door-to-Door Canvassing Experiments.” *Journal of Politics* 65:1083-96.

Hoel, Jessica, Jeffrey Parker and Jon Rivenburg. 2004. “Peer Effects: Do First-Year Classmastes, Roommates, and Dormmates Affect Students’ Academic Success.” Paper presented at the Higher Education Data Sharing Consortium Winter Conference, January 14, 2005. Retrieved from http://www.reed.edu/~hoelj/HEDS.pdf.

Huckfeldt, Robert and John Sprague. 1995. *Citizens, Politics, and Social Communication: Information and Influence in an Election Campaign*. New York: Cambridge University Press.

Iyengar, Shanto and Donald R. Kinder. 1987. *News that Matters*. Chicago: University of Chicago Press.

Jennings, M. Kent and Richard G. Niemi. 1974. *The Political Character of Adolescence: The Influence of Families and Schools*. Princeton: Princeton University Press.

Katz Lawrence F., Jeffrey R. Kling and Jeffrey B. Liebman. 2001. Moving to opportunity in Boston: early results of a randomized housing mobility study. *Quarterly Journal of Economics* 116:607–54.

Key, Vladimer Orlando. 1949. *Southern Politics in State and Nation*. New York: Alfred A. Knopf.

Mackie, Diane M., Leila T. Worth, and Arlene G. Asuncion. 1990. “Processing of Persuasive In-Group Messages.” *Journal of Personality and Social Psychology* 58:812-822.

Manski, Charles F. 1993. “Identification of Endogenous Social Effects: The Reflection Problem.” *Review of Economic Studies* 60:531-42.

Marmaros, David, and Bruce Sacerdote. 2002. “Peer and Social Networks in Job Search.” *European Economic Review* 46: 870-879.

McEwan, Patrick J. and Kristin A. Soderberg. 2004. “Peer effects at a single-sex college.” Unpublished manuscript, Wellesley College.

Mutz, Diana C. 1998. *Impersonal Influence*. New York: Cambridge University Press.

Newcomb, Theodore M. and G. Svehla. 1937. “Intra-family Relationships in Attitude,” *Sociometry*, 1:180-205.

Nickerson, David W. 2005. *Measuring Interpersonal Influence*. Doctoral dissertation: Yale University.

Oliver. J. Eric. 2001. *Democracy in Suburbia*. Princeton, NJ: Princeton University Press.

Sacerdote, Bruce I. 2001. "Peer Effects With Random Assignment." *Quarterly Journal of Economics*, 116:681-704.

Sacerdote, Bruce. 2003. "Peer Effects with Random Assignment: Results from Dartmouth Roommates." Quarterly Journal of Economics, p. 681-704.

Stinebrickner, Todd R. and Ralph Stinebricker. 2001. “Peer Effects Among Students from Disadvantaged Backgrounds.” Working paper: The University of Western Ontario.

Wilson, John, and Marc Musick. 1997.Who Cares? Toward an Integrated Theory of Volunteer Work.” *American Sociological Review* 62:694-713.

Winston, Gordon C., and David J. Zimmerman. 2003. “Peer Effects in Higher Education” (Working Paper no. 9501). Cambridge, MA: National Bureau of Economic Research. Retried from http://www.nber.org/papers/w9501.

Table 1 OLS Regression Results Modeling the Peer Effects of Roommates and Suitemates
I'm sure this table is eloquent, but I have trouble figuring out what it is saying (without a lot of effort). Is it okay or do we want to make it more accessible?

Variable	Discuss Politics	Number of days a week you read a newspaper	Number of days a week watch television	Drinking Age	Favorable / Unfavorable Pot Smokers	Legalize Marijuana	Favorable / Unfavorable Homosexuals	Homosexual Teachers
Constant	-0.06	-0.10	-0.33	-0.59***	0.23	0.09	0.20	-0.53***
	(0.14)	(0.26)	(0.47)	(0.13)	(0.14)	(0.14)	(0.12)	(0.12)
Pretest Answer	**0.57*****	0.25***	0.34***	0.37***	0.70***	0.66***	0.56***	0.36***
	(0.07)	(0.03)	(0.04)	(0.06)	(0.05)	(0.06)	(0.05)	(0.04)
Pretest Roommate	0.18**	0.12***	0.18***	0.15**	0.19***	0.14**	0.24***	0.10**
	(0.07)	(0.04)	(0.05)	(0.06)	(0.06)	(0.07)	(0.05)	(0.05)
Pretest Suitemate	0.21*	0.13**	0.15*	0.20**	0.24**	0.16	0.16	0.14
	(0.12)	(0.06)	(0.08)	(0.10)	(0.10)	(0.12)	(0.10)	(0.09)
Female	-0.24	-0.12	-0.33	0.30*	0.05	-0.02	0.46**	-0.38***
	(0.18)	(0.16)	(0.20)	(0.17)	(0.20)	(0.18)	(0.19)	(0.12)
N	204	230	230	202	197	198	204	202
Adj-R-sq	0.30	0.23	0.26	0.20	**0.51**	**0.38**	**0.51**	**0.36**

The dependent variable is the post-test response of the student.
Unless otherwise stated, questions were on a 5-point agree/disagree scale.
Favor/unfavorable ratings are on a 7-point scale.
implies p<0.1; ** implies p<0.05; *** implies p<0.01

Table 2 Effect of Peer-to-Peer Mobilization

	Internet Volunteers			Vote for America (Rhode Island)		
	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3
Discussion (α)	0.048	-0.022	-0.046	-0.030	-0.038	-0.032
	(0.064)	(0.051)	(0.051)	(0.048)	(0.035)	(0.033)
Prior Voter History		Yes	Yes		Yes	Yes
Member of Major Party			0.190**			0.275**
			(0.045)			(0.033)
Captain Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Constant	0.412**	0.268*	0.219	0.224	0.078	-0.062
	(0.149)	(0.121)	(0.118)	(0.205)	(0.152)	(0.143)
N	374	374	374	481	481	481
Adj-Rsq	0.13	0.43	0.47	0.06	0.48	0.55

Dependent variable is voter turnout in the 2002 general election.
Coefficients are derived from two-stage least squares using assignment as an instrument for contact.
Numbers in parentheses represent standard errors.
* implies p<0.05 and ** implies p<0.01

Pooled together, these experiments estimate the mobilization to be –3.6 percentage points with a standard error of 2.8 percentage points. However, the upper bound of the 95% confidence interval for the pooled result is a meager 1.7 percentage points.

Table 3 Contagion Effect of Voting within the Household (2SLS)

	Model 1	Model 2	Model 3
Contagion (α)	**0.614****	**0.564***	**0.587****
	(0.233)	(0.220)	(0.217)
Voted 2001		0.165**	0.155**
		(0.040)	(0.037)
Voted 2000 General		0.041	.041
		(0.034)	(0.035)
Voted 2000 Primary		0.134*	0.119*
		(0.053)	(0.049)
Voted 1999		0.128**	0.120**
		(0.045)	(0.042)
Voted 1998 General		-0.001	-0.014
		(0.030)	(0.028)
Age			-0.002
			(0.004)
Age-Squared			0.000
			(0.000)
Female			0.020
			(0.027)
Denver	0.059	0.044	0.019
	(0.055)	(0.046)	(0.046)
Constant	0.072	-0.113**	-0.087
	(0.054)	(0.033)	(0.101)
N	955	955	955
Adjusted R sq	0.45	0.56	0.55

Dependent variable = Voter turnout in 2002 Primary of person in household who did not answer the door.
Endogenous variable = Voter turnout in 2002 Primary of person in household who did answer the door.
Instrument = Assignment to treatment or placebo condition.
Voter History is for person who did not answer the door (same as the dependent variable).
Numbers in parentheses represent standard errors.
* implies $p < 0.05$ and ** implies $p < 0.01$