8. Political trust in experimental designs

Rick K. Wilson and Catherine C. Eckel

INTRODUCTION

In her 1997 Presidential Address to the American Political Science Association, Elinor Ostrom pointed to the importance of trust for solving collective action problems and she called for laboratory experiments to tease out why trust could act as a mechanism for facilitating such solutions (1998). Around the same time the Russell Sage Foundation developed an initiative to study trust, focusing on its role in enabling and supporting a wide variety of relationships, from personal interactions to those involving economic exchange and business dealings. Several important publications came out of this initiative. Levi and Stoker (2000) reviewed the literature on trust and trustworthiness up to that point and called on the community to utilize experimental methods to address the strategic aspects of trust. Hardin (2002) elaborated on many of the issues of trust that had been ignored by the formal theory community. Finally the workshop and resulting volume edited by Ostrom and Walker (2003) brought together an array of scholars from many different disciplines to address different aspects of trust. These included not only social science theorists and experimentalists, but also evolutionary psychologists and animal behavior experts. While these efforts helped push the use of experiments to study trust, as Nannestad (2008) argued in part, the lessons learned were unclear. Our goal here is to illustrate how experiments are informing unresolved questions about political trust.

As every chapter in this Handbook notes, trust and trustworthiness are critical for understanding social behavior. Many see political trust as a necessary foundation for stable political institutions and for the formation of social capital and civic engagement (Putnam, 1993, 2000; Stolle, 1998). Others focus on the extent to which particular aspects of political institutions – such as transparency, appropriate levels of monitoring, and credible threats of sanctions – can foster political trust and cooperation (Arrow, 1974; Fenno, 1978; Bianco, 1994; Lupia and McCubbins, 1998; Sztompka, 1999; Knight, 2001; Ostrom, 2003). Finally trust is seen as a social lubricant that reduces the cost of exchange, whether in reaching political compromise (Fenno, 1978; Bianco, 1994) or in daily market and non-market exchange interactions (Arrow, 1974). While political trust both affects and is affected by political institutions, the causal relationships between levels of trust and effectiveness of institutions have been difficult to disentangle from observational data. Increasingly, experiments, both in the lab and in the field, focus on the causal underpinnings of political trust.

Political scientists have long been concerned with the concept of political trust and its relation to political institutions. This tradition extends back to Almond and Verba (1963), who find a strong correlation between citizen trust levels and the existence of democratic institutions. Subsequent work has examined the nature and causal direction of this relationship using non-experimental data. One difficulty in assessing research in this area is that the concept of trust is not clear-cut, with different researchers addressing

or utilizing very different measures of trust. Like others in the Handbook we begin with Pippa Norris's (see also Chapter 2) nested dimensions of political support to provide focus for the concept of political trust. In particular, we show how experiments lead to greater understanding of public confidence in institutional principles, political institutions, and in office-holders. Each is addressed through the discussion of three different experimental designs: a survey experiment, a public goods experiment, and a trust experiment. We then turn to measurement issues that arise when focusing on behavior. Finally, we conclude with open questions about political trust that experiments can address.

THE ROLE OF EXPERIMENTS IN THE STUDY OF TRUST

Research using experimental methods is not a panacea for addressing all aspects of trust and trustworthiness: an experiment is merely one of the many methodological tools available to a researcher. An experiment involves two features: control by the experimenter and random assignment over the treatments. By necessity, to gain control over naturally occurring variation, the experimenter abstracts the setting in which the subject acts. The aim is to focus on a particular (predicted) relationship. The experimenter then randomly assigns subjects to a treatment in order to see if the predicted effect occurs in the presence of the treatment and disappears when the treatment is removed.

Experiments are especially valuable for addressing questions of causality and this is where they shine. Likewise, experiments are useful for exploring and pinpointing mechanisms that researchers think are critical for understanding complex social phenomena. In a different vein, sometimes multiple theories can be used to explain a given phenomenon. Experiments can be valuable in distinguishing between competing theories, allowing researchers to focus on those that are the most promising. Experiments need not always test theories – sometimes they can produce new facts that are valuable for theory building. Observing an unexpected relationship in the laboratory is eye-opening and often leads to new theoretical insights as well as additional experiments designed to test the new theory (and to confirm that the observation is not simply anomalous). Finally, the incentivized tasks developed for experiments are often then used as measures, and as we will demonstrate, canonical experiments can be valuable measurement tools.

Careful control, a focus on mechanism, and randomization allow experimenters to make specific causal claims. But, this 'internal validity' is achieved at a price. McDermott (2011) elaborates on the trade-offs between internal and external validity. When used in conjunction with other methods, social scientists can make clear causal claims that can be generalized to settings outside the lab, or outside a specific field context. Given the complex claims made about political trust, experiments have an important role to play in understanding the concept and its impact in a variety of specific settings.

THREE CANONICAL EXPERIMENTAL DESIGNS

In their survey of the literature Levi and Stoker (2000) lamented the fact that there were few laboratory experiments focused on political trust. A huge body of experimental work on trust has appeared since then. However, that research focuses on social trust

in dyads rather than trust by citizens toward political institutions (see the survey by Wilson and Eckel, 2011). With respect to political trust, little experimental work has surfaced. Sigelman et al. (1992) is an early example that uses an experiment to causally sort between trust in political institutions and partisanship. Substantial research, particularly in political psychology, focuses on the antecedents to political trust. This includes procedural fairness, competency, and performance. Grimes (see also Chapter 16) summarizes the empirical and experimental literature in this arena. Other exceptions include the laboratory experiments by Mutz and Reeves (2005), Ahn et al. (2008), Dal Bo et al. (2010), Boudreau (2012), and Grimmelikhuisen (2012). What surprises us is the limited experimental work explicitly addressing the concept of political trust.

We focus on three types of experimental designs that have been used by trust researchers. These designs are not exhaustive, but each is chosen because it tackles a specific dimension of the trust relationship as elaborated by Norris (see also Chapter 2). We find these designs to be promising for future work and we focus on specific studies to highlight the promise of experiments for the study of political trust.

Design One: Survey Experiments

Increasingly, political scientists rely on survey experiments to test for attitude change, the effects of framing, or to use priming to clarify cognitive differences among subjects. As with all experiments, survey experiments involve the researcher designing treatments and randomly assigning respondents to the treatments. Typically, a survey experiment can involve larger samples of the population than are feasible with other types of experiments, and can be designed to allow for planned subgroup analysis. Also by using a large sample, many treatments can be simultaneously carried out. Mutz (2011) provides a very good overview of the advantages (and cautions) associated with survey experiments.

This type of experiment is likely to be most familiar to scholars studying attitudes about political trust, and therefore we take this as a natural starting point. A good deal of research has been devoted to understanding the correlation between citizen trust and political institutions. In this vein Hetherington and Kam (2013) provide a useful example of a survey experiment. Starting from a concern for the extent and impact of climate change, Hetherington and Kam ask whether cues from highly trusted governmental institutions can be used to change public opinion in the direction of greater support for climate change policy. In framing the question in this manner they engage Norris's (see also Chapter 2) fourth dimension of political trust: public confidence in political institutions. They propose a causal mechanism between trust in a government institution and attitudes held by citizens and then use an experiment to test for that mechanism.

Hetherington and Kam (2013) draw on a rich literature on public opinion change in political science, and argue that effective cue-givers are 'trusted, credible and liked.' Equally important they draw on the finding that people attend more carefully to 'counter-stereotypical' information (Rahn, 1993) – that is, information that contradicts their stereotype about the institution. Their key question is whether a 'trusted' source providing 'counter-stereotypical' information leads to a change in policy attitudes.

Measuring attitudinal change is difficult. The typical approach is to use a cross-section of the population, measure attitudes, collect covariates, and then estimate the effect of beliefs about organizations on attitudes, controlling for the covariates. A more rigorous

design, but one that is difficult to execute, is to use panel data employing repeated measures of key attitudes. Problematic, however, is that panel data are often collected far apart in time, and many factors may vary in that interval. Drawing an inference about any change in attitudes then relies on an assumption that intervening events have had no systematic, confounding effects. This is where an experiment shines. The researchers can manipulate both the source and type of information that respondents receive. With random assignment to treatments, differences in attitudes between the different treatments can be ascribed with confidence to the treatments.

Hetherington and Kam (2013) begin with data from the General Social Survey (GSS) indicating that the US military is the most liked and trusted American political institution. At the same time, it is not viewed as holding liberal views about the environment. Accordingly, they ask whether a strong environmental message originating from the US military can change public opinion about climate change. The military should be a useful cue-giver because it is trusted. Moreover, if the cues that it gives are counter-stereotypical, in that they contradict the belief that the military is not environmentally concerned, then it should be more effective in changing public opinion.

In their experiment Hetherington and Kam (2013) vary two factors. First is the source of information, which comes either from the federal government or from the military. The information statement common to both states that the organization (federal government or military) is turning to alternative (non-fossil) fuel sources for strategic and long-term cost reasons. Following the statement respondents are asked their views about whether global warming is an important problem.

The second factor in the experiment varies the information about the target of an appropriation to increase energy efficiency. The statement says that a \$2 billion appropriation will be given to the military or the Environmental Protection Agency (EPA) to increase energy efficiency. Respondents are then asked whether they agree or disagree with giving the appropriation to that entity. This generates a 2×2 factorial design in which the source of message (federal government or military) is manipulated and in which the target (military or EPA) is also manipulated. They find that counter-stereotypical endorsements have a very strong effect. Those who trust the military, but would not expect the military to be a strong supporter of environmental innovation, are the most persuaded to re-evaluate their views about climate change and appropriations. Not surprisingly the effects are largest for Republicans – a group that has high trust in the military, but is suspicious of environmental spending. Using a survey experiment, Hetherington and Kam have sufficient numbers of respondents to include these subgroup effects.¹

Survey experiments are increasingly used in order to isolate causal effects. Surprisingly only a handful of such experiments have been used to investigate trust in government (see, for example, Brooks and Geer, 2007; Jacobs and Matthews, 2012; Craig and Rippere, 2013). When using a survey experiment in the field, researchers are able to draw a causal inference with some external validity. For example, survey experiments are easy to carry out on a representative population sample. The major disadvantage with survey experiments in the field is that one loses the tight control that is afforded by the laboratory. Unlike in the laboratory some respondents may elect not to respond, leaving missing data that might be systematically tied to the treatment. This introduces potential selection bias into the inference. In addition, some respondents may not have been attentive to the treatment, thereby biasing downward the treatment effect. Despite potential pitfalls,

survey experiments are useful for testing conjectures about the trust relationship. One particularly thorny problem for observational studies has been to disentangle the causal effect of trust in political institutions. A survey experiment helps clarify the relationship.

Design Two: Public Goods Games

A second type of design typically takes place in the laboratory and relies on the canonical public goods game. In essence the game is an N-person prisoner's dilemma (PD) game, and is typically implemented in the laboratory using a voluntary contribution mechanism (VCM), where subjects voluntarily contribute to the provision of a public good. In its simplest form, each member of a group of N players must choose to allocate a fixed initial endowment either to a group account (a public good) or to a private fund (that is, 'defect' in PD language). Each unit invested in the public good produces a certain amount of the public good, and each person in the group receives the same return from this public good, which is called the marginal per capita return (MPCR). For concreteness, suppose that a unit allocated to the group account produces 0.5 units of the public good. Then each person in the group, whether she contributed or not, receives earnings of 0.5 units. The dominant strategy is to keep one's contribution while sharing in the contributions by others: this is equivalent to free riding.

In a one-shot version of the game, theory predicts that everyone will keep their endowment and no public good will be provided. In the repeated version of the game, where the same group interacts over several rounds, the incentives remain the same. This version of the VCM has the same set of subjects play ten or more rounds with feedback about group contributions in each round. No one is expected to contribute and everyone is expected to free ride. Scores of experiments show that between 40 and 60 percent of subjects contribute in the one-shot game (Dawes, 1980; Dawes et al., 1988). In the repeated version of the experiment contributions begin around the same rate in the first period and then decline toward zero as the experiment continues (Ledyard, 1995).

How does the structure of this game relate to trust? Orbell et al. (1984; Orbell and Dawes, 1993) conjectured that the choice of contributing was tied to beliefs about others in the group. These beliefs included the possibility of being taken for a sucker and the possibility that one could trust others to contribute. These latter beliefs are most closely related to *generalized* trust – whether those in the group can be trusted to contribute to the public good.²

Beginning with the idea that cooperation is a measure of who can be trusted, Strimling et al. (2013) tackle one of the quintessential questions in the trust literature: can institutions alter levels of trust? In framing the question in this manner they engage Norris's (see also Chapter 2) second dimension of political trust – whether there is support for the core institutional principles. Strimling et al. (2013) assert a causal direction: weak institutions will lower levels of generalized trust.

In their design, subjects first are randomly assigned to groups of four or five. Each is given an endowment of 10 in each of 20 periods, and they can choose to divide the amount between their private account and the group account in any way. Whatever is placed in the group account is multiplied by 1.6 (for groups of four) or 2.0 (for groups of five), and then divided evenly among all players in the group (thus the MPCR is 0.4). This standard VCM is played for six periods. Subjects are then sorted into two groups: those who contribute

more and those who contribute less than the median. As other research has shown, this effectively sorts individuals into two types – those who tend to be cooperative and those who tend to be uncooperative (Gunnthorsdottir et al., 2007). Blocking by type, subjects are then randomly assigned into one of two different institutions to again play the VCM.

In both institutions a smallest acceptable level of contribution to the common pool is exogenously imposed. At a personal cost, a subject can choose to monitor the contributions of another group member. Doing so means that another member of the group is randomly chosen and audited to see if that person complied with the minimum contribution rule. If a non-complier is found, then the monitor receives a reward, which exceeds her cost of monitoring, and that money is withdrawn from the common pool. The non-complier is assessed a fine. Two types of institutions are used: the ‘weak institution’ (their term) has a low acceptable level of contribution (1 unit) and a low fine (2 units). By comparison the ‘strong institution’ (their term) has a high acceptable level of contribution (8 units) and a high fine (9 units). Every third round the group can vote whether to alter the institution. They can vote to decrease by 1 unit the minimum acceptable level of contribution or fine; do nothing; or increase by one unit the minimum acceptable level of contribution or fine. Although the acceptable level of contribution is bounded by 0 and 9 (the latter to allow some monitoring), the level of the fine is unbounded at the upper end.

The key results are easy to see. First, those in the strong institution tend to retain that institution, even though they could have voted to make it less onerous. Second, there are differences between cooperative and uncooperative types. Among the former, most groups stay with the high acceptable level of contribution. Among the latter, a majority of the groups quickly moved to undermine the acceptable level of contribution, but then moved back to increasing the acceptable level. Third, those in a weak institution react differently. Cooperative types step up the acceptable level of contribution until the weak institution is now a strong institution. This is not the case for non-cooperative types. The weak institution does not give them insight into the gains that can be had from full cooperation and those groups do the worst in terms of payoffs.

With respect to fines, there are similar differences. In the strong institution, non-cooperative types decrease fines at first, and then raise them back to levels that are slightly higher than where they started. Cooperative types steadily increase the fines over the course of 20 rounds such that they go from 9 units to almost 12.5 units on average. In the weak institution, cooperative types steadily increase fines such that they are almost 8 units by period 20. This is not the case for the non-cooperative types in the weak institution. Again, these individuals simply cannot get the group to a socially efficient point.

Strimling et al. (2013) conclude that low-trust groups (non-cooperators) struggle in weak institutions. However, strong institutions can build trust for non-cooperators. This lends credence to claims made by others that political institutions can build trust in society. At the same time, it is clear that high-trust groups (cooperators) can cope with a variety of institutions. Those individuals are less vulnerable to weak institutions.

Here is an instance in which the careful design of an experiment allows researchers to investigate the causal direction between trust and institutions. It may be that contributions to the public goods game are weak proxies for generalized trust. Nonetheless the findings are intriguing for understanding the role that institutions play for increasing cooperation. Because the researchers manipulated the type of institution and the subjects were randomly assigned to the type of institution, the results are clear. Of course this is not

the last word on the effect of institutions on trust. For example, Bohnet et al. (2001) find that institutions designed to promote dyadic trust appear to ‘crowd out’ trust. It may be that the coercive elements of ‘strong institutions’ means that trust is unnecessary. If the institution perfectly monitors all citizens and perfectly enforces compliance, then what role does trust play? The effect of institutions on political trust remains an open, but important, research question.

Design Three: The Investment Game

The third experimental design is the trust game, also known as the investment game. This design is most closely aligned to the *particularistic* (dyadic) trust relationship (Hardin, 2002). Developed by Berg et al. (1995), the experiment involves a pair of individuals playing a sequential game in which both players receive an initial endowment. The first player, the trustor, can take any part of her endowment and send it to the second player, the trustee. Whatever is sent is tripled by the experimenter and given to the second player. The second player then decides how much, if anything, to return to the first player. Here the trustor has made herself vulnerable to the second player, trusting that the trustee will return something. What the first player sends is considered to be a measure of *trust*. What is returned is considered to be a measure of *trustworthiness*. A large number of studies have been carried out using this design; see the review by Wilson and Eckel (2011). The usual finding is that a substantial proportion of the population is willing to put their trust in others. As well, the level of trustworthiness is usually sufficiently high that it pays, on average, to trust the second mover.

We report on results from one of our own experiments. In doing so we engage Norris’s (see also Chapter 2) fifth dimension of political trust – confidence in officeholders. Rather than using the investment game as an experiment to test a specific treatment, we use it to *measure* the level of trust in government: we assess trust in public officials by citizens. Note that this is different from asking subjects about their attitudes toward public officials; instead we have a direct behavioral measure of political trust in which subjects risk a non-trivial amount of money by trusting a public official. We study two small towns in Texas that are matched on size, similarities in their economy, similarities in their political structure, but varying in their risk from natural disasters. As part of the study we measure the levels of trust between citizens and public officials.³

Two prior studies have used public officials in experimental settings. This, of course, is difficult to do in the laboratory since public officials have constraints on their time. Enemark et al. (2013) use a group of local politicians in Zambia and Butler and Kousser (2013) use a sample of US state legislators. Both find that public officials are willing to trust one another (and in the latter case do so at higher rates than a comparable student sample). While these findings are interesting, they do not address the question of whether citizens trust public officials. In our study we play the dyadic trust game between citizens and public officials. In each community, a sample of individual citizens and a sample of local officials are recruited.⁴ The study is an example of a laboratory experiment in the field. It has characteristics of the laboratory (control by the experimenter), but targets non-students as subjects, and is carried out in settings that are convenient for the subjects.

Two variations of the trust game are discussed here.⁵ In the first game both first and second movers are citizens. In the second game the first mover is a citizen and the second

mover is a public official.⁶ Both players start with a \$30 endowment. The first mover chooses how much of his \$30 he would like to send to the other player. Unlike the standard trust game, subjects are told they can only send amounts in \$10 increments (\$0, \$10, \$20, \$30). Whatever is sent is tripled, and the second mover then decides how much, out of the tripled money plus her own original \$30, to return to the first player. For the second mover, any amount (including zero or her full budget) can be sent back to the first mover.

Prior to making a choice, subjects are told that they would be randomly assigned to be either a first mover or a second mover and that they will be randomly paired with another person in the room. Because no one knew the role they will play, they need to make a choice in both roles. If the game is selected for payment, subjects are randomly assigned to a role and randomly paired.

Subjects are then asked to make a decision as the first mover. We also asked subjects how much they expect to be returned from the second mover. Immediately following that decision, subjects are asked to make decisions as second movers. In this part of the experiment we use the 'strategy method', so for any combination of money that was sent (and tripled) subjects are asked how much they intend to return, thereby specifying their full strategy as a second mover. This meant that subjects had four boxes with \$30, \$60, \$90, and \$120 and they are asked, for whatever amount is sent and tripled, what they will return. Subjects are also asked how much they expect to be sent. Subjects receive no feedback about their counterpart's choice until the end of the experiment.

The second game is almost identical to the first. The primary difference involves the roles of players. Citizens are told they will be first movers in this game and they need to decide how much of their \$30 endowment they wish to send to second movers. Public officials are assigned as second movers and both citizens and public officials are reminded that they will be matched with one another if the game is drawn. Otherwise all aspects of the second game are the same as the first.

Do citizens trust fellow citizens more than their public officials? The answer is yes. We use what was sent in the investment game to be a measure of trust. In game one the target of trust is another citizen while in the second game the target is a public official. Figure 8.1 plots the percentages sent for each of the possible categories. Because of differences between the two cities, Panels A and B separately represent the cities. Pooling all of the data, we note that the distributions are different – subjects tend to send less to the public officials (Wilcoxon signed-rank test, $z = 3.28$, $p < 0.001$). Both panels make this point. The percentage sending money to public officials shifts toward zero, indicating that fellow citizens are more trusted. A striking point about the figure is how few subjects sent nothing and how many subjects sent \$20 or \$30. In fact, in City 1, just over 45 percent sent their entire endowment to fellow citizens.

It may be that citizens had accurate expectations about their local public officials. We asked citizens, in their role as first mover, how much they expect to be reciprocated conditional on what they sent. They did this for both games. At the same time, we collected information about what subjects said they would return for each amount that could be sent to them (if they were assigned to be a second mover). We took the average expected by a citizen given what they sent and subtracted that from what individuals said they would return, on average, for that amount. So, in City 1, people who sent nothing (failed to trust) expected that the other player would return \$7.50. On average, citizens indicated that they would send back \$4.33. What is interesting is that these citizens are dipping into their own endowment to send something to their counterpart. Taking the same example, citizens

Note: Standard errors are plotted on the figure.

Figure 8.1 Percentage of citizens sending \$0, \$10, \$20 or \$30 to their counterpart (citizen, public official)

sending nothing expected less returned from their public officials (\$2.78) and on average public officials indicated they would send back \$0.96. Across the board subjects expected less from their public officials than their fellow citizens (a within-subjects pairwise test: $t = 2.78$, $df = 133$, $p < 0.01$).

Figure 8.2 gives a sense of whether subject's expectations were correct. Again we find differences between the two cities. In City 1, expectations are always higher than what either citizens or public officials will return. For City 2, however, public officials are likely to return more than expected by subjects. Subjects have low expectations for the public officials (lower than their fellow citizens), yet except for the case when nothing is sent, the public officials always indicate they will send more.

We have managed to do something unusual – pairing citizens and public officials in the trust game. The only difference between the two games is the identity of the counterpart. We have three findings. First, citizens are willing to trust their public officials. Second, citizens' willingness to trust is driven by expectations. Third, those expectations and their accuracy vary by context. Unfortunately, we have only two different cities (contexts). While both are similar along many dimensions, we cannot pinpoint why citizens have such low expectations in City 2. It may be that the public officials feel citizen pressure or it may be that public officials in this disaster-prone city feel compelled to be more trustworthy. We cannot isolate these effects with this limited sample.

Notes:

The difference is the average of expectations conditional on what a first mover sent minus the average that could have been returned for the amount sent.
Standard errors are plotted on the figure.

Figure 8.2 Difference in expected and actual returns by counterpart (citizens or public officials)

In our study we primarily use the investment game as a measurement tool. As such we want to know whether it is correlated with subjects' willingness to comply with post-disaster instructions from local officials. At the same time, we want to know whether this costly measure of behavior can be proxied with standard attitudinal trust questions. Table 8.1 provides the bivariate correlations with a number of attitudinal items that we collected as part of the study. The first column correlates the amount sent in the investment game to fellow citizens. The second column does the same for what is sent to public officials. The first item is one of our key variables of interest – whether subjects will comply with local

Table 8.1 *Bivariate correlations with amount sent to citizens and public officials*

	Trust in Citizens	Trust in Public Officials
When a disaster happens how likely are you to follow directions given by the local police and fire officials? (5-point scale)	0.075	0.078
How much do you trust your neighbors? (4-point scale)	0.108 †	−0.015
How much do you trust city government? (4-point scale)	−0.001	0.036
Generally speaking, would you say that most people can be trusted, or that you need to be very careful in dealing with people? (11-point scale)	−0.131*	0.071

Notes:

All questionnaire items are set in the same direction with higher values indicating greater trust.

† $0.5 < p < 0.10$; * $p < 0.05$.

officials. The correlation between the behavioral measure and the attitude is positive but weak. Using a variety of attitudinal trust measures we find there are equally weak correlations with the behavioral measures. The standard GSS question about trust is positively correlated with trust in citizens and is significant. However, the relationship is quite weak. This leaves us with a conundrum: what does the investment game measure? In the next section we point out that the measurement issues we find are widespread.

MEASUREMENT CONCERNS

A key concern for those carrying out studies of any kind is to measure accurately. An experiment allows control over the environment and the subjects have an incentive to behave according to their true preferences. Both factors should lead to precision in measurement as compared with survey questions. However, as we show above, there is little correlation between attitudes measured with survey items and incentivized choices. What does the literature have to say about why this might occur?

The comparisons between attitudinal and behavioral measures almost exclusively focus on generalized trust. Yet the same comparisons should be undertaken with respect to political trust. An early study on generalized trust by Glaeser et al. (2000) used a variety of survey and behavioral measures in a laboratory setting. Subjects filled out an attitudinal survey, engaged in hypothetical choices and played the incentivized ‘investment game’. As with our study, the standard GSS attitudinal item was included (‘Generally speaking, would you say that most people can be trusted or that you can’t be too careful in dealing with people’) as well as a number of other items commonly used to measure attitudes toward trust. They also collected information about trusting behavior, such as whether someone locks their door or lends money to friends. They found (like us) that the attitudinal survey items were uncorrelated with the experimental measure of trust, but that the self-reported behavior is correlated with behavior in the

experiment. At best the GSS items may capture ‘trustworthiness’. But behavior in the game was a pretty good predictor of self-reported trusting behavior. That attitudes and behavior can be uncorrelated is not surprising – this has been a long-standing conclusion reached by psychologists (Ajzen and Fishbein, 1977; Eagly and Chaiken, 1993). This study has led those carrying out experiments to focus more on behavior rather than attitudes.

Others have expressed concern about what the standard GSS item measures. For example, Miller and Mitamura (2003) contend that the standard GSS item taps ‘caution’ rather than trust. The standard wording they argue has two parts and each part is conceptually distinct. ‘Obviously it is possible for a risk-averse person to feel that people in general are trustworthy, but still to be inclined to be careful in dealing with others’ (Miller and Mitamura, 2003, pp. 63–4). Reeskens and Hooghe (2008) ask whether the standard GSS item can be cross-culturally compared. Like Miller and Mitamura (2003) they demonstrate that the structure of the item is problematic. For example, the item shows high variability even among the same respondents (measurement error), the wording is problematic, and it is not clear to the respondent who the appropriate target (‘other people’) might be. In the end they conclude that there is considerable instability in the measure.

Mixed results concerning what is meant by the behavioral and attitudinal measures of generalized trust have caught the eye of experimentalists. In a thorough study of the tie between attitudinal and behavioral measures of trust, Capra et al. (2008) find that there is considerable variability in which attitudinal measures are correlated with which behavioral measures. Their results are consistent with those noted above. The usual GSS item predicts both trustworthiness and trust in the investment game, if one controls for other-regarding preferences. The attitudinal statement that ‘I am trustworthy’ is *negatively* correlated with reciprocity in the trust game, confirming the old adage of never trusting someone who says ‘trust me’. Similar findings are also reported by Fehr (2009). There is little consensus about how these attitudinal and behavioral measures fit together.

In the end, measurement problems abound for the study of generalized trust. We have every reason to believe that the same may be true for political trust. The concept is complicated and unlikely to be captured by any single measure. Different methods can be used to chip away at the concept and different methods will focus on unique dimensions of political trust. Laboratory experiments have promise as a useful tool for establishing the reliability and validity of measures of key dimensions of political trust.

UNRESOLVED QUESTIONS

We have provided three examples of canonical experiments used to study aspects of political trust. Each touches on a different dimension of Norris’s (see also Chapter 2) concept of political trust. We argue that trust experiments are particularly useful for isolating mechanisms that create political trust. Experiments allow a researcher to control the stimulus that a subject is given and to control the subject’s environment. Rigorous control comes at a price. It may be difficult to generalize to a broader set of contexts from findings uncovered by experiments in a specific lab setting. We do not see this as a problem for research using experiments. We advocate using experiments in conjunction with observational and archival data. No single method has a monopoly on truth. Multiple methods

force researchers to ask the same question from different perspectives and, in doing so, leads to greater confidence in results when the findings converge. From our perspective there are three unresolved questions that experiments can shed light on. First, it is important to establish the causal relationship between political trust and political institutions. The question is whether trust precedes effective institutions or vice versa. The contours of this question are now being tackled, as noted by several of the studies mentioned above. Here experiments are particularly well suited to establishing the causal relationship.

Second, we are puzzled by the extent to which attitudinal trust items relate to trust behaviors in natural settings. The inconsistency in measurement should give both experimenters and observational researchers pause. The concept of political trust has many dimensions and it is likely that few of our 'workhorse' measures are sufficient to explain all of them. Experiments can be a useful tool for developing and testing measurements.

Third, we do not know the extent to which strong institutions crowd out trust. It may be that strong institutions are a substitute for the norms of trust and trustworthiness. If so, strong institutions that can successfully monitor and sanction individuals may make trust unnecessary. But such institutions may be overly coercive. Instead political trust may be critical for placing a check on coercive political institutions.

NOTES

1. A second experiment corrects for threats to inference embedded in the first design, which combined both the expertise and endorsement of the military. Here they use a 2×4 factorial design (the details of which the reader will be spared). They find, consistent with their first study, that it is not the endorsement alone that is necessary, but also the expertise held by the political institution. This leads to the conclusion that trust in institutions is shored up by expertise.
2. Such experiments are not confined to the laboratory (see, for example, English, 2012). Buchan et al. (2009) use the public goods game to assess whether exposure to globalization leads to greater trust. They use a within-subject design to measure contributions not only in groups that include people in their own locality, but also in groups that contain other people scattered across the globe. As subjects are increasingly exposed to globalization they are more likely to contribute to global public goods.
3. This study is part of a larger project that is designed to explore whether citizens and public officials share similar preferences and concerns about natural disasters (Donahue et al., 2014).
4. The public officials included elected officials (mayors, county judges/commissioners, and city council members) and appointed civil servants (including city managers, emergency managers, first responders, and others involved in city planning).
5. Subjects completed eight different games in total, with one selected randomly for payment. The remaining six games are not explicitly relevant to the research discussed here. For example, individuals also completed incentivized measures of risk and time preference, as well as charitable giving games and risk-sharing games.
6. Public officials also participated in a version of the first game. They were told they were participating with other public officials and they would be matched with another public official.

REFERENCES

- Ahn, T.K., R.M. Isaac and T.C. Salmon (2008), 'Endogenous group formation', *Journal of Public Economic Theory*, **10** (2), 171–94.
- Ajzen, I. and M. Fishbein (1977), 'Attitude–behavior relations: A theoretical analysis and review of empirical research', *Psychological Bulletin*, **LXXXIV**, 999–1018.
- Almond, G.A. and S. Verba (1963), *The Civic Culture*, Princeton, NJ: Princeton University Press.
- Arrow, K.J. (1974), *The Limits to Organizations*, New York: Norton.

- Berg, J.E., J.W. Dickhaut and K. McCabe (1995), 'Trust, reciprocity, and social history', *Games and Economic Behavior*, **10** (1), 122–42.
- Bianco, W.T. (1994), *Trust: Representatives and Constituents*, Ann Arbor, MI: University of Michigan Press.
- Bohnet, I., B.S. Frey and S. Huck (2001), 'More order with less law: On contract enforcement, trust, and crowding', *American Political Science Review*, **95** (1), 131–44.
- Boudreau, C. (2012), 'Greater than the sum of their parts? When combinations of institutions improve citizens' decisions', *Journal of Theoretical Politics*, **24** (1), 90–109.
- Brooks, D.J. and J.G. Geer (2007), 'Beyond negativity: The effects of incivility on the electorate', *American Journal of Political Science*, **51** (1), 1–16.
- Buchan, N., G. Grimalda and R.K. Wilson et al. (2009), 'Globalization and human cooperation', *Proceedings of the National Academy of Science*, **106** (11), 4138–42.
- Butler, D.M. and T. Kousser (2013), 'How do public goods providers play public goods games?', paper presented at the 2013 Annual State Politics & Policy Conference and the 2013 Annual Meeting of the American Political Science Association.
- Capra, M.C., K. Lanier and S. Meer (2008), *Attitudinal and Behavioral Measures of Trust: A New Comparison*, Atlanta, GA: Emory University.
- Craig, S.C. and P.S. Rippere (2013), *Political Trust and Negative Campaigns: Two Tests of the Figure-Ground Hypothesis*, Gainesville, FL: University of Florida.
- Dal Bo, P., A. Foster and L. Putterman (2010), 'Institutions and behavior: Experimental evidence on the effects of democracy', *American Economic Review*, **10** (5), 2205–29.
- Dawes, R.M. (1980), 'Social dilemmas', *Annual Review of Psychology*, **31**, 169–93.
- Dawes, R.M., A.J.C. Vandekragt and J.M. Orbell (1988), 'Not me or thee but we – the importance of group identity in eliciting cooperation in dilemma situations – experimental manipulations', *Acta Psychologica*, **68** (1–3), 83–97.
- Donahue, A.K., C.C. Eckel and R.K. Wilson (2014), 'Ready or not? How citizens and public officials perceive risk and preparedness', *American Review of Public Administration*, **44** (4), 89–111.
- Eagly, A.H. and S. Chaiken (1993), *Psychology of Attitudes*, Fort Worth, TX: Harcourt Brace Jovanovich.
- Enemark, D.P., C.C. Gibson, M.D. McCubbins and B.A. Zimmerman (2013), 'The effect of holding office on the behavior of politicians', accessed 26 July 2016 at http://polisci2.ucsd.edu/denemark/papers/Enemark_et_al_PoliticianExperiments.pdf.
- English, W. (2012), 'Demystifying trust: Experimental evidence from Cambodia and Thailand', *Journal of Theoretical Politics*, **24** (2), 172–209.
- Fehr, E. (2009), 'On the economics and biology of trust', *Journal of the European Economic Association*, **7** (2–3), 235–66.
- Fenno, R.F., Jr. (1978), *Home Style: House Members in Their Districts*, Boston, MA: Little, Brown and Company.
- Glaeser, E.L., D.I. Laibson, J.A. Scheinkman and C.L. Soutter (2000), 'Measuring trust', *The Quarterly Journal of Economics*, **115** (3), 811–46.
- Grimmelikhuijsen, S. (2012), 'Linking transparency, knowledge and citizen trust in government: An experiment', *International Review of Administrative Sciences*, **78** (1), 50–73.
- Gunnthorsdottir, A., D. Houser and K. McCabe (2007), 'Disposition, history and contributions in public goods experiments', *Journal of Economic Behavior and Organization*, **62** (2), 304–15.
- Hardin, R. (2002), *Trust and Trustworthiness*, New York: Russell Sage Foundation.
- Hetherington, M.J. and C. Kam (2013), *Institutional Cue-Giving and Persuasion: Enlisting the Military as Environmental Protector*, Nashville, TN: Vanderbilt University.
- Jacobs, A.M. and J.S. Matthews (2012), 'Why do citizens discount the future? Public opinion and the timing of policy consequences', *British Journal of Political Science*, **42** (4), 903–35.
- Knight, J. (2001), 'Social norms and the rule of law: Fostering trust in a socially diverse society', in K.S. Cook (ed.), *Trust in Society*, New York: Russell Sage Foundation, pp. 354–73.
- Levi, M. and L. Stoker (2000), 'Political trust and trustworthiness', *Annual Review of Political Science*, **3** (1), 475–507.
- Lupia, A. and M.D. McCubbins (1998), *The Democratic Dilemma: Can Citizens Learn What They Need to Know? Political Economy of Institutions and Decisions*, Cambridge, MA: Harvard University Press.
- McDermott, R. (2011), 'Internal and external validity', in J.N. Druckman, D.P. Green, J.H. Kuklinski and A. Lupia (eds), *Cambridge Handbook of Experimental Political Science*, Cambridge, UK: Cambridge University Press, pp. 27–40.
- Miller, A.S. and T. Mitamura (2003), 'Are surveys on trust trustworthy?', *Social Psychology Quarterly*, **66** (1), 62–70.
- Mutz, D.C. (2011), *Population-Based Survey Experiments*, Princeton, NJ: Princeton University Press.
- Mutz, D.C. and B. Reeves (2005), 'The new videomalaise: Effects of televised incivility on political trust', *American Political Science Review*, **99** (1), 1–15.

- Nannestad, P. (2008), 'What have we learned about generalized trust, if anything?', *Annual Review of Political Science*, **11** (1), 413–36.
- Orbell, J.M. and R.M. Dawes (1993), 'Social-welfare, cooperators advantage, and the option of not playing the game', *American Sociological Review*, **58** (6), 787–800.
- Orbell, J.M., P. Schwartz-Shea and R.T. Simmons (1984), 'Do cooperators exit more readily than defectors', *American Political Science Review*, **78** (1), 147–62.
- Ostrom, E. (1998), 'A behavioral approach to the rational choice theory of collective action', *American Political Science Review*, **92** (1), 1–22.
- Ostrom, E. (2003), 'Toward a behavioral theory linking trust, reciprocity, and reputation', in E. Ostrom and J.M. Walker (eds), *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, New York: Russell Sage Foundation, pp. 19–79.
- Ostrom, E. and J.M. Walker (eds) (2003), *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, New York: Russell Sage Foundation.
- Putnam, R.D. (1993), *Making Democracy Work*, Princeton, NJ: Princeton University Press.
- Putnam, R.D. (2000), *Bowling Alone: The Collapse and Revival of American Community*, New York: Simon & Schuster.
- Rahn, W. (1993), 'The role of partisan stereotypes in information processing about political candidates', *American Journal of Political Science*, **37** (2), 472–96.
- Reeskens, T. and M. Hooghe (2008), 'Cross-cultural measurement equivalence of generalized trust: Evidence from the European Social Survey (2002 and 2004)', *Social Indicators Research*, **85** (3), 515–32.
- Sigelman, L., C.K. Sigelman and B.J. Walkosz (1992), 'The public and the paradox of leadership: An experimental analysis', *American Journal of Political Science*, **36** (2), 366–85.
- Stolle, D. (1998), 'Bowling together, bowling alone: The development of generalized trust in voluntary associations', *Political Psychology*, **19** (3), 497–525.
- Strimling, P., S.I. Lindberg and M. Ehn et al. (2013), *Building and Maintaining Strong Institutions: An Experimental Approach*, Gothenburg: University of Gothenburg.
- Sztompka, P. (1999), *Trust: A Sociological Theory*, Cambridge, UK: Cambridge University Press.
- Wilson, R.K. and C.C. Eckel (2011), 'Trust and social exchange', in J.N. Druckman, D.P. Green, J.H. Kuklinski and A. Lupia (eds), *The Handbook of Experimental Political Science*, Boston, MA: Cambridge University Press, pp. 243–57.