



Handbook of Trust Research

Edited by **Reinhard Bachmann** and **Akbar Zaheer**



3 Can groups be trusted? An experimental study of trust in collective entities

Bill McEvily, Roberto A. Weber, Cristina Bicchieri and Violet T. Ho

Introduction

Trust is the topic of a considerable amount of recent research in the social sciences. This trend is particularly noteworthy in the economics, organizational and strategy literatures, where trust is considered extremely important for many kinds of interaction. For instance, several economists argue that trust is an essential 'lubricant' without which even the simplest forms of economic exchange can not occur (Arrow 1974).¹ Trust increases the efficiency of exchange by reducing the expectation of opportunistic behavior and consequently lowering associated transaction costs (Bromiley and Cummings 1995; John 1984; McEvily and Zaheer, chapter 16, this volume). Strategy researchers suggest that trust is a strategic resource that has the potential to provide a source of sustained competitive advantage (Barney and Hansen 1995), while other organizational researchers conceptualize trust as a governance form that provides a framework to guide and direct the organization and coordination of economic activity (Bradach and Eccles 1989; McEvily et al. 2003; Powell 1990).

Incorporating the concept of trust into economic, strategic and organizational theories clearly holds the potential of producing far-reaching implications for our understanding of exchange, competition and behavior in economic and organizational settings. By focusing on the motives and intentions of economic actors, this line of research promises to explicitly investigate and sharpen the core assumptions upon which theory is based. At the same time, however, integrating the concept of trust into existing theory poses a number of challenges. Chief among these is the question of how to extend – or whether it is reasonable to extend – an individual-level construct such as trust to more aggregate levels of analysis.

To a large extent placing trust in individuals and placing trust in collective entities (e.g. groups, organizations, industries, institutions, etc.) are used interchangeably in the literature and without specific consideration for whether differences in the object of trust are meaningful or appropriate. For example, transaction costs economics proposes that 'human agents are given to opportunism' (Williamson 1985, p. 64), but that *firms* must safeguard their transactions against the threat of such untrustworthy behavior. As a result there is ambiguity about the object of trust that is most relevant to minimizing transaction costs of exchange – the individual agent or the partner organization. This raises several questions. First and foremost, does trust exist at different levels of analysis (individual versus collective)? If so, is trust across levels related and does trust at one level influence trust at another? For instance, if one trusts the individual agent with whom one deals, then is one necessarily more inclined to trust that person's organization? Or is it even possible to trust a collective entity, independent of the trust one has for its individual

members? And, perhaps most importantly, does trust at different levels of analysis affect economic behaviors in different ways?

Surprisingly, these fundamental questions have received relatively little research attention despite the widespread application of trust to economic and organizational relationships involving collective entities. A related stream of research in sociology focuses on 'system' (Luhmann 1979; Giddens 1990) and 'institutional-based' (Zucker 1986) trust. These forms of trust refer to abstract structures that shape expectations through generalized rules of behavior. System and institutional-based trust create commonly accepted background assumptions, and thereby lower the inherent risk of trusting a counterpart (Bachmann 2001; 2003). Similar to our notion of trust in collective entities, system and institutional-based trust are 'impersonal' forms of trust that are not based on familiarity with a specific individual. At the same time, our view of trust in a collective entity presupposes that the individual is involved in a direct relationship (economic exchange in this study) with the collective entity, whereas system and institutional-based trust primarily refer to the economic framework within which the relationship is embedded. For instance, in the context of buyer-supplier interfirm exchanges, trust in a collective entity would be a purchasing manager's trust in the buyer organization while system and institutional-based trust would focus on, for example, the legal system of contract enforcement, regulatory agencies, third-party brokers, etc.

Although the work on system and institutional-based trust is conceptually related to our focus on trust in a collective entity, there is virtually no empirical work validating the existence of trust in collective entities. A notable exception is a study by Zaheer et al. (1998) that explores, using survey responses of boundary-spanning agents in buyer-supplier interfirm exchanges, the relationship between trust in a specific individual dealt with and trust in that individual's organization, referred to as interpersonal and inter-organizational trust respectively. Zaheer et al. (1998) find that interpersonal and inter-organizational trust are related, but distinct; economic agents discriminate between the two types of trust, but also view trust at different levels of analysis as strongly related. Although their study provides evidence consistent with the concept of trust in a collective entity, it does not definitively rule out the possibility that trust in a collective entity (i.e. inter-organizational trust) is merely an artifact of trust in the individual members of the collective entity (i.e. interpersonal trust). In particular, it may be the case that trust in a collective entity is simply a function of trust in its individual members. If so, trust in a collective entity is reducible to trust in individuals and does not exist as a separate concept. This suggests that validating the existence of trust in a collective entity requires evidence that is not specifically tied to trust in individuals and raises the question of whether there is some aspect of trust in a collective entity that exists apart from trust in the members of a collective entity.

The purpose of this chapter is to explore the relationship between trust in individuals and trust in collective entities, and the relationship of these two types of trust to economic behavior. We argue that trust can be meaningfully applied to economic transactions involving both individual and collective entities, and that trust at these two levels is related. We propose that economic actors form perceptions about the trustworthiness of collective entities based on exchanges conducted with individual members of the collectivity. This trust in the collective then becomes transferable to other individuals within the collectivity and serves as a proxy for individual trust where detailed knowledge of

individual members of the collectivity is limited or absent (McEvily et al. 2003; Stewart 2003). This is true even when there is no reason to believe that these other individuals are likely to exhibit similar trust-related properties. Thus trust in the collective entity is used as a heuristic for trust in individuals and is extended to transactions with other members of the collectivity, even those that are unknown and about which there is little or no information for determining trustworthiness. In this way, membership in a collectivity or group can be taken to signal trustworthiness (Kramer et al. 1996), without each member having to exhibit his or her trustworthiness directly to all other parties in an economic exchange.

Based on the above ideas, we address the following research question: *Is an individual's perception of a counterpart's trustworthiness affected by the counterpart's membership in a group and by the past actions of others in that group, even in a situation where membership in the group conveys no meaningful information about trustworthiness?*

To explore this question we conducted a laboratory experiment where the outcomes of economic exchanges were influenced by the degree to which subjects trusted their counterparts and the extent to which those counterparts actually upheld the trust that was placed in them. Our research builds on previous experiments using the 'trust' (or 'investment') game (e.g. Berg et al. 1995). We extend this paradigm by including a treatment in which we embed exchange within the context of minimally defined collective entities. In this treatment, transactions among individual economic actors are no longer isolated events, but rather are linked through individuals' membership in collective entities.

Conceptualizing trust in a collective entity

Trust is an inherently complex concept (Corazzini 1977) that has been studied from a number of different disciplinary perspectives. As a result, a wide variety of definitions exist. Despite the heterogeneity in conceptualizations, there are a number of common elements unifying the many different usages of trust. In particular, there is widespread agreement that *trust is the willingness to be vulnerable based on the positive expectation of the intentions or behavior of others* (Mayer et al. 1995; Rousseau et al. 1998). Moreover, for trust to arise, interdependence and uncertainty are necessary conditions. Interdependence means that the interest of one party cannot be fulfilled without reliance on another party. Uncertainty means that the possibility of experiencing negative outcomes by relying on another party requires taking a 'leap of faith' (Lewis and Weigert 1985). If another's intentions could be ascertained with complete certainty, trust would not be needed. Accordingly, trust is the choice to make oneself vulnerable under the conditions of interdependence and uncertainty.

Consistent with the broader literature on trust we refer to the extent to which one individual trusts another individual with whom she deals as *trust in an individual*. The degree to which a sales representative trusts the specific purchasing manager with whom she deals is an example of trust in an individual. In contrast, *trust in a collective entity* represents the extent of trust that an individual places in a collectivity with which she deals. Trust in a collective entity exists between an individual on the one hand and a collection of individuals on the other. A sales representative's trust for the buyer organization that she transacts with is an example of trust in a collective entity. The distinction between trust in an individual and trust in a collective entity is based on the object of trust. Whereas the source of trust resides in individuals for both, the object of trust differs. Rather than being directed at a specific individual, the referent of trust in a collective entity is an aggregate

social system comprising a number of individuals. The placing of trust in a collective entity, rather than a specific individual, is consistent with definitions of trust that emphasize 'confidence in or reliance on some quality or attribute of a *person or thing*' (*Oxford English Dictionary*, emphasis added).²

While the conceptual distinction between trust in an individual and trust in a collective entity is fairly well established, empirical evidence substantiating the distinction is largely absent. As noted previously, in order to distinguish the two forms of trust it is important to separate the trust an individual places in a collective entity from the trust that an individual places in individual members of the collective entity. For instance, an individual may claim to trust a certain organization, but may only be referring to the trustworthiness of specific individuals in the organization or of the general population from which the organization draws its membership. Similarly, one may generally believe that people are trustworthy, and therefore most organizations comprising ordinary people are also trustworthy. Or, one may feel that a group is trustworthy because of familiarity with all of the members of the group and their trustworthiness. We argue that these are not instances of trust in a collective entity, since there is no separate attribution of trustworthiness to the organization as an entity in itself. Instead, we suggest that trust in a collective entity *can be more cleanly* distinguished from trust in individuals by studying a behavioral manifestation of trust that is clearly separable across, and distinctly attributable to, the two types of trust.

To disentangle trust in individuals and trust in collective entities, we designed a laboratory experiment that allowed us to directly explore the possibility that trust may exist for groups, independent of the trust for the individuals in those groups. In the experiments, we created a very basic form of collective using a variant of the well-known 'minimal group paradigm' (Tajfel et al. 1971). Research on minimal groups shows that there is a discontinuity between individual and group behavior: people tend to behave differently when confronting another individual or a group, or when they themselves act as group members. What is even more striking is that this discontinuity occurs even when the group is created on the basis of an inconsequential criterion, and group membership is anonymous. We thus expected to find a difference in behavior when subjects were faced with 'unlabeled' individuals as opposed to members of a designated group. In the absence of previous interactions with a specific member of a group, we were interested in exploring whether an initial experience with an anonymous member of the same group translates into a stereotypical judgment of the whole group, on some chosen dimension.

In particular, we explored the extent to which subjects were likely to exhibit trust in another subject based on experience with a previous member of that subject's group. Specifically, participants in our experiment played the trust game twice. Our focus is on how the actions of the first counterpart affected decisions when playing the game with the second counterpart. We are especially interested in whether membership of the two counterparts in the same 'minimal group' makes this effect stronger. We find a modest effect, in both magnitude and significance. However, the presence of a positive result is compelling, even given the modest significance, since we used the weakest possible form of group identity in our experiments. More generally, the finding is important because it provides a clear and carefully controlled demonstration that lends support to the existence of trust in collective entities.

The trust game

As its name suggests, the trust game creates a situation where one player must decide whether to trust another, and this other must then decide whether to honor or abuse this trust. Specifically, Player 1 is given some initial wealth allocation of which she must decide how much to 'trust' to Player 2. Player 2 can be thought of as an agent of Player 1 who has the ability to turn this trusted amount into an even greater sum. Therefore, the amount received by Player 2 is some multiple of the amount trusted to Player 2 by Player 1. After receiving this amount, Player 2 must decide how much, if any, of the total amount received to return to Player 1.

This game models several situations in which the attractiveness to one party of a welfare-increasing investment hinges on the trustworthiness of another. For instance, consider a situation where the owner of a small firm has to decide how much training to provide an employee. This training is costly for the owner of the firm, but can yield greater profits for both the employee and the firm, provided the employee remains with the firm after the training. Once the owner decides how much to commit to training and the training actually takes place, the employee then decides how long to remain with the firm. Assuming that the employee can realize greater profit by leaving to go to another firm once the training is received, the problem is exactly the one modeled by the trust game. Player 1 (the owner) decides how much of some allocation to commit to Player 2 (the employee), who then decides whether to honor this trust (remain with the firm, in which case both employee and owner receive a better payoff than if there had been no training) or abuse this trust (leave the firm immediately after training, yielding the highest payoff to the employee but the lowest to the owner).

The game can also be described formally. In the continuous version of the game, Player 1 is given some amount $W > 0$, which she can divide between one amount she keeps for herself and one she trusts to Player 2. Label the amount she trusts to Player 2 as x , with $0 \leq x \leq W$. The amount x is then multiplied by a constant, $r > 1$, so that the second player receives the greater amount rx . Player 2 must then decide what proportion, k , of rx to return to Player 1, keeping the rest, $(1-k)rx$, for himself. Assuming that Player 2 also receives some fixed sum c (which might be zero), the following are the payoffs for the game:

$$\begin{aligned} \text{Payoff to Player 1:} & \quad \pi_1 = (W - x) + xrk = W + (kr - 1)x \\ \text{Payoff to Player 2:} & \quad \pi_2 = c + (1 - k)rx \end{aligned}$$

Player 2 moves second and the choice of k does not affect x , which has already been determined. Therefore, as long as Player 2 is maximizing his monetary payoff, he will select k equal to zero and keep the entire amount rx . Knowing this, Player 1 should always keep the entire amount W and set x equal to zero. Thus, in the unique subgame-perfect Nash equilibrium to the game, $x = k = 0$, $\pi_1 = W$, and $\pi_2 = c$.

The game is interesting, however, because trust on the part of Player 1 can lead to an outcome that Pareto-dominates (i.e. is more efficient than) this equilibrium. This is true for any outcome in which x is greater than zero and kr is greater than one, meaning that Player 1 invests a positive amount and receives more than that amount back from Player 2.

Several laboratory experiments studied the trust game. In the first example of such a study, Berg et al. (1995) used the trust game to determine whether or not trusting behavior can be found when social enforcement is not possible. In their experiments subjects

played the game in an environment where the usual self-interested motivations assumed by economists to lead to trusting behavior were eliminated. Subjects played the game only once and under complete (double-blind) anonymity. In spite of this anonymity and lack of repetition, only two of the 32 subjects in the role of Player 1 sent \$0. On the other hand, five subjects sent the entire amount of \$10. The average amount sent was \$5.16 and the average amount returned was \$4.66, indicating that sending money led to slight losses on average for Player 1.³

Taken together, the experiments by Berg et al. and others using different variations of the trust game – with varying payoffs and parameters – show some consistent results, even across cultures (e.g. Van Huyck et al. 1995; Güth et al. 1997; Snijders and Keren 1998; Buchan et al. in press; Ashraf et al. 2005). First, the subgame-perfect equilibrium prediction is rarely observed. Most subjects in the role of Player 1 send a positive amount to Player 2. On the other hand, most subjects who sent money as Player 1 did not send the full amount W . A second main finding is that while many subjects in the role of Player 2 returned a positive amount to Player 1, the returns tended to be slightly less than the original investment on average. Therefore, while subjects in general exhibited trusting behavior, this trust was often repaid, but usually not sufficiently to prevent it from being costly. Consequently, in our experiments, which use a variation of the trust game, we expect a significant amount of trusting behavior. However, our attention is primarily on whether trusting behavior is influenced by past experience with counterparts who belong to the same group.

One previous study explored the connection between group boundaries and trust, although with an entirely different focus than ours. Buchan et al. (2002) used random assignment to divide subjects into two groups (Proposers and Responders) and then used the trust game to measure the extent to which subjects in the first group exhibited trust for subjects in the second group. The treatment variable was the nature of the relationship between Proposers and Responders. In a Direct condition, a Proposer sent money to a Responder who then sent money back to the same Proposer. In a Group condition, Proposer A sent money to Responder A while Proposer B sent money to Responder B, and Responder A then sent money back to Proposer B while Responder B sent money back to Proposer A. In this condition, reciprocity was indirect, but two Proposers and two Responders were mutually linked by their actions. Finally, in a Society condition, Proposer A sent money to Responder B, who sent money back to a randomly selected Proposer C. In this condition, reciprocity was indirect and links between Proposers and Responders were much more distant than in the Group condition. The results in all three conditions revealed significant amounts of trust and reciprocation, though both of these decreased as the interaction between Proposers and Responders became less direct. Buchan et al.'s experiments demonstrate that trust exists even when it involves indirect reciprocation between members of randomly determined groups, but that this trust (measured by the amount sent by Proposers) is less the more indirect the relationship.

The study by Buchan et al. is relevant for our experiment since it shows that subjects exhibit trusting behavior even when the object of this trust is not directly responsible for reciprocating it. One interpretation of their results is that, even with groups determined by an entirely random process, subjects are willing to trust counterparts when someone else in this counterpart's group must reciprocate this trust. In this case the object of trust seems to be the group rather than a specific individual.⁴ While the experiments do not

constitute a direct test of trust in a collective entity, the results are consistent with the notion that subjects can trust a group rather than an individual.

Our experiments differ from those of Buchan et al. in that we *directly* explore trust in a collective entity. In particular, we focus on whether a subject's past history of dealing with one member of a group influences that subject's propensity to trust another member of the same group, beyond the information that such history provides about the trustworthiness of the second individual.

Experimental design

Our experiment tests whether subjects assigned minimal group labels use these labels to draw inferences about the trustworthiness of other individuals. In our experiment, subjects play the trust game twice against two subjects randomly selected from the population of other participants. Our treatment variable is the relationship between these two other subjects. In the Control condition, they are simply referred to as two other subjects of the opposite role (Player 1 or Player 2), which was randomly determined at the beginning of the experiment. In the Group condition, these two other subjects are members of the same 'minimal group' that was determined by responses to an unrelated question. We are particularly interested in how subjects respond to the outcome of the first game, when playing the second game. Our hypothesis is that subjects in the Group condition will be more influenced by what their first counterpart did than those in the Control. Therefore, our experiment is primarily intended to test whether perceptions of trustworthiness are transferred more readily across individuals who are in the same group than across individuals with no such group label.

Subjects in our experiment played two rounds of the following discrete version of the trust game:

- Player 1 was given an allocation of \$4 at the beginning of the game.
- Player 1 then chose an amount to send to Player 2. This amount was \$0, \$2 or \$4.
- Player 2 received an amount equal to four times the amount sent by Player 1.
- Player 2 then decided whether to return to Player 1 either \$0 or half of the amount received.

Note that this is the same as the trust game discussed in the previous section, with $W = \$4$, $x \in \{\$0, \$2, \$4\}$, $r = 4$, $c = \$0$, and $k \in \{0, \frac{1}{2}\}$. Therefore, the payoffs to Player 1 and Player 2, respectively, were:

$$\begin{aligned}\pi_1 &= 4 + (4k - 1)x \\ \pi_2 &= (1 - k)4x.\end{aligned}$$

As in other versions of the trust game, the unique subgame-perfect Nash equilibrium is for Player 1 to send \$0 and for Player 2 to return \$0 for any amount received, leaving Player 1 with \$4 and Player 2 with \$0. However, this equilibrium outcome is Pareto-dominated by the outcome in which Player 1 sends \$4 and Player 2 returns half, leaving both players with \$8.

Each of the sessions in our experiment consisted of 10–20 subjects recruited from a distribution list of students at Carnegie Mellon and the University of Pittsburgh. At the beginning of the session, subjects were divided into two groups (explained in more detail below). Each subject then played the game twice, in the same role, with two randomly

selected subjects from the other group. Subjects did not know the identity of the other subjects with whom they were playing the game.

In each play of the game, actions were made and recorded using a choice sheet. At the beginning of the game, Player 1 circled on the choice sheet how much he or she wanted to send to Player 2. The sheet was then collected, the choice recorded, and the sheet was given to a Player 2. This Player 2 then circled his or her choice of how much to send back to Player 1.⁵ The sheets were then collected, the choices recorded, and the sheet was given back to Player 1 who could observe the outcome of the game. Players also had record sheets on which they recorded what happened in each of the two games.

The only difference between the two treatments was in how the groups were determined and in the labels used to refer to the two roles.

- In the Control condition, subjects were randomly assigned participant numbers at the beginning of the experiment. They were then told that odd participant numbers corresponded to the role of Player 1 and that even participant numbers corresponded to the role of Player 2. Subsequently, the two roles were referred to as 'Player 1' and 'Player 2'.
- In the Group condition, subjects were also randomly assigned participant numbers, but these were not used to determine the roles. Instead, subjects were asked to make a guess about the number of days it would rain the following year in San Francisco. A median split of these guesses was then used to divide the subjects into two groups: High Guessers and Low Guessers. High Guessers played the role of Player 1, while Low Guessers played the role of Player 2. Subsequently, all reference to the two roles was made using the terms 'High Guessers' and 'Low Guessers'.

Note that this is a very weak group manipulation. In one condition, the roles are simply determined by a guess about something unrelated to the game. There was no other difference between the two treatments. In both treatments, subjects who were in the role of Player 1 were seated on one side of the room while subjects in the role of Player 2 were seated on the other. Subjects were visible to each other during the experiment, but did not know with whom they were matched.⁶ At the end of the experiment, subjects completed a questionnaire measuring their general propensity to trust (Rotter, 1967).

We conducted 12 sessions (six in the Control condition and six in the Group condition), using a total of 174 Carnegie Mellon and University of Pittsburgh graduate and undergraduate students (80 in Control and 94 in Group). Subjects were recruited from a large pool of potentially interested participants via an e-mail announcement that provided little information on the details of the experiment. The sessions were conducted between September 2000 and May 2001.

Results

Our main hypothesis is that the interaction between experience (what happens in the first round) and experimental treatment affects the amount sent in the second round by Player 1. Specifically, the presence of trust in a collective entity implies that subjects in the Group condition will be more influenced by experience than subjects in the Control. We begin our analysis by exploring the aggregate data for other patterns of behavior related to the group manipulation.⁷

Aggregate behavior

Table 3.1 presents the total amounts sent by subjects in the role of Player 1 by condition. The aggregate choices by subjects do not differ greatly by condition. There are slightly more Player 1s who sent \$4 in the Control (60 percent) than in the Group condition (48 percent), but this difference is not significant. Moreover, almost twice as many subjects in the Group condition (32 percent) than in the Control (18 percent) initially sent \$2. In fact, while about 80 percent of subjects in the role of Player 1 in both conditions sent some amount of money in the first round, a larger proportion of those sending some money sent \$4 in the Control condition (25 of 32, 78 percent) than in the Group condition (22 of 37, 59 percent). This difference in amount sent among those who sent money is marginally significant in a Fisher Exact test ($p = 0.08$). However, this pattern is reversed – but is not significant – when we look at the choices in Round 2. Specifically, in the second round we see that a smaller proportion of those sending some money sent \$4 in the Control condition (23 of 29, 79 percent) than in the Group condition (23 of 28, 82 percent).⁸ Note also that in both treatments the frequency of players sending \$0 increased between Rounds 1 and 2, and that this increase was greater in the Group condition (from 21 to 40 percent) than in the Control condition (from 20 to 28 percent). Overall, among subjects in the role of Player 1, there are slight differences in behavior between the two conditions when looking at the aggregate data. In particular, of those sending some money, the distributions of amounts sent differ between the treatment and control conditions.

Table 3.2 reports the behavior of subjects in the role of Player 2 by condition and round. Each entry in the table gives – for each possible amount sent – what proportion of Player 2s returned one-half of the amount received, resulting in an improvement for Player 1 over the initial allocation. The remaining subjects all returned \$0, resulting in a

Table 3.1 *Frequencies of amounts sent by Player 1*

Condition	Amount sent	Round 1	Round 2	Total
Control	\$0	8 (20%)	11 (28%)	19 (24%)
	\$2	7 (18%)	6 (15%)	13 (16%)
	\$4	25 (63%)	23 (58%)	48 (60%)
Group	\$0	10 (21%)	19 (40%)	29 (31%)
	\$2	15 (32%)	5 (11%)	20 (21%)
	\$4	22 (47%)	23 (49%)	45 (48%)

Table 3.2 *Percentage of Player 2s returning half by offer*

Condition	Amount sent	Round 1	Round 2	Total
Control	\$2	4/7 (57%)	1/6 (17%)	5/13 (38%)
	\$4	13/25 (52%)	12/23 (52%)	25/48 (52%)
	Total	17/32 (53%)	13/29 (45%)	30/61 (49%)
Group	\$2	3/15 (20%)	3/5 (60%)	6/20 (30%)
	\$4	11/22 (50%)	12/23 (52%)	23/45 (51%)
	Total	14/37 (38%)	15/28 (54%)	29/65 (45%)

loss for Player 1. Cases where Player 2 received \$0 meant there was no subsequent choice and are therefore not included in the table.

Again, when we look at the aggregate data we see small differences between the two treatments. Note first that in Round 1, the number of subjects who returned half is greater in the Control (53 percent) than in the Group condition (38 percent), but this difference is not significant. This difference is largest for subjects who were sent \$2. In the Control condition, 4 of 7 (57 percent) such subjects returned one-half; in the Group condition, only 3 of 15 (20 percent) such subjects did so. This difference, however, is not significant in a Fisher Exact test ($p = 0.11$). When pooling across rounds and amounts sent, we see that Player 2s in the Control condition were only very slightly more likely to return half (49 percent) than those in the Group condition (45 percent).

Individual behavior and trust in a collective entity

When a subject in the role of Player 1 sends either \$2 or \$4 to a Player 2, the outcome is one of two possibilities: either half the multiplied amount is returned or nothing is returned. In one case, Player 1 is better off – relative to the initial allocation – for having sent an amount greater than \$0, and in the other Player 1 is worse off. Therefore, we can think of these as situations where initial trust in Round 1 is either ‘honored’ or ‘abused’. The main aim of this chapter is to explore what happens in Round 2 when trust is either honored or abused in Round 1 and, in particular, whether subjects in the Group condition are more affected by these events than those in the Control sessions. Such behavior is directly relevant for testing our main prediction that trust in a collective entity is separable from trust in the individual members of that collectivity. The results in Tables 3.1 and 3.2, however, do not allow us to test this hypothesis because the results are presented at an aggregate level across rounds. Instead it is necessary to examining subjects’ experiences across rounds and within conditions, which is how the data are organized in Table 3.3.

Table 3.3 Choices in Round 2 by Player 1 contingent on outcomes in Round 1

Sent in Round 1	Returned in Round 1	Trust honored or abused	Sent in Round 2	Control	Group
\$0	N/A	No information	\$0	5 (13%)	7 (15%)
			\$2	3 (8%)	1 (2%)
			\$4	0 (0%)	2 (4%)
\$2	\$0/\$8	Trust abused	\$0	3 (8%)	8 (17%)
			\$2	0 (0%)	2 (4%)
			\$4	0 (0%)	2 (4%)
\$2	\$4/\$8	Trust honored	\$0	0 (0%)	0 (0%)
			\$2	1 (3%)	0 (0%)
			\$4	3 (8%)	3 (6%)
\$4	\$0/\$16	Trust abused	\$0	3 (8%)	4 (9%)
			\$2	0 (0%)	2 (4%)
			\$4	9 (23%)	5 (11%)
\$4	\$8/\$16	Trust honored	\$0	0 (0%)	0 (0%)
			\$2	2 (5%)	0 (0%)
			\$4	11 (28%)	11 (23%)
Total				40	40

Table 3.3 presents, for all possible outcomes of Round 1, the subsequent Round 2 choices of subjects in the role of Player 1 in each condition. The first two columns in the table present the possible outcomes in the first round. The next column classifies these outcomes into three possible categories from Player 1's point of view: no information (if \$0 was sent and no action of Player 2 was observed), trust abused (if either \$2 or \$4 was sent and \$0 was returned), and trust honored (if either \$2 or \$4 was sent and one-half of the multiplied amount was returned). The fourth column presents the possible amounts sent in Round 2 by a Player 1, and the last two columns give the number of subjects in each condition who sent that amount after observing the outcome described in the first three columns.

The results in the table reveal greater sensitivity to prior outcomes in the Group condition than in the Control condition. For instance, of those subjects in the role of Player 1 who sent \$4 in Round 1 and received back \$0 (i.e. trust was abused), 9 of 12 subjects (75 percent) in the Control condition again sent \$4. In the Group treatment, however, only 5 of 11 such subjects (45 percent) again sent \$4. Similarly, of the subjects who sent \$4 in Round 1 and received \$8 back (i.e. trust was honored), all of the 11 subjects in the Group condition again sent \$4, but a smaller fraction (9 of 11; 85 percent) did so in the Control group. However, neither of these differences alone is significant in a Fisher Exact test. Still, these results suggest a greater sensitivity to first-round results on the part of subjects in the Group condition than the Control condition.

A direct test of our hypothesis involves looking at how subjects in the role of Player 1 react when their initial trust is either abused or honored. To demonstrate the existence of trust in a collective entity, we need to show that subjects in the Group condition whose trust is abused (honored) in Round 1 are likely to send less (more) in Round 2 than subjects in the Control condition whose trust is abused (honored). Table 3.4 presents the relevant results. Specifically, the Round 2 choices of subjects in the role of Player 1 are given, by condition and outcome in Round 1. Using Table 3.4, we can see whether behavior in the two conditions differs in the way we predicted, and in a way consistent with subjects

Table 3.4 Choices of Player 1 in Round 2 by Round 1 outcomes and condition

Amount sent in Round 2	Trust honored in Round 1		Total
	Control	Group	
\$0	0 (0%)	0 (0%)	0
\$2	3 (18%)	0 (0%)	3
\$4	14 (82%)	14 (100%)	28
Total	17	14	31

Amount sent in Round 2	Trust abused in Round 1		Total
	Control	Group	
\$0	6 (40%)	12 (52%)	18
\$2	0 (0%)	4 (17%)	4
\$4	9 (60%)	7 (30%)	16
Total	15	23	38

displaying trust in a collective entity, even with the minimal form of groups created in our experiments.

As the top part of Table 3.4 indicates, in the Group condition, all 14 subjects (100 percent) whose trust was honored subsequently sent \$4, while in the Control 14 of 17 subjects (82 percent) did so and the other 3 only sent \$2 (18 percent). While the direction of this difference – that subjects in the Group condition whose trust is honored are slightly more likely to send \$4 in the next round – is consistent with our hypothesis, deviations by three subjects is insufficient to produce a significant difference.

Additional, and more compelling, direct support for our hypothesis can be found in the bottom part of Table 3.4. Here, we explore the behavior of subjects in the role of Player 1 who had their trust abused in the first round (they sent some amount of money and received \$0 in return). There is a clear difference in the pattern of choices between the two conditions. In the Control, a majority of subjects (60 percent) whose trust was abused still sent \$4 in the next round. In the Group condition, however, only 30 percent of such subjects did so, and a majority of subjects (52 percent) sent \$0. The difference between the distributions of actions in the two conditions is significant at the $p < 0.1$ level in a chi-square test ($\chi^2(2) = 4.78$).⁹ Therefore, we see a significant difference in behavior consistent with our hypothesis of trust in a collective entity: subjects in the Group condition whose trust is abused in Round 1 were significantly more likely to react negatively in Round 2 than those in the Control.¹⁰

We can also explore our hypothesis using regression analysis to determine the effect of first-round experience on behavior in the second round. Specifically, we used ordered logit estimation to explore how the amount sent in Round 2 is affected by the experimental treatment variable, Round 1 history, and an interaction between treatment and history.¹¹ The results of this estimation are reported in Table 3.5.¹² The first three independent

Table 3.5 Ordered logit regression of amount sent in Round 2

Dependent variable: amount sent in Round 2	(1)	(2)	(3)	(4) ^a	(5) ^a
Trust abused (Rd 1)		-2.670*** (0.696)	-2.599*** (0.729)	-1.518** (0.827)	-1.150* (0.884)
Group condition		-0.286 (0.576)	-0.315 (0.598)	0.955 (1.211)	1.055 (1.221)
Trust abused x Group				-1.901* (1.385)	-2.206* (1.422)
Dispositional trust	0.644* (0.395)		0.154 (0.579)		0.419 (0.565)
Gender (Male = 1)	-0.533 (0.432)		-0.141 (0.630)		-0.594 (0.624)
Obs	87	69	69	69	69
Log likelihood	-81.76	-48.88	-48.83	-50.55	-49.89
Pseudo R-squared	0.027	0.185	0.186	0.182	0.192

Notes:

Standard errors are in parentheses.

^a Value of one observation of dependent variable changed (see note 12).

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; all one-tailed.

variables are binary variables indicating whether trust was abused in Round 1, whether the subject was in the Group condition, and the interaction between the two. To test robustness, we also include a gender dummy variable and a construct from questionnaire responses measuring an individual's general propensity to trust.

As the results indicate, if trust was abused, subjects sent significantly less than if it was not. However, the negative effect of trust being abused on amount sent in the subsequent round is even stronger for subjects in the Group condition (at least twice as big). The coefficient on the interaction term ($p < 0.09$ in model 4 and $p < 0.06$ in model 5, one-tailed) is statistically significant providing further support for our hypothesis.

Conclusion

The above experiment and analysis reveal evidence of trust in a collective entity. Subjects in the Group condition show a greater reaction to previous experience, particularly when this experience is negative. While the effect is not small, it is only weakly statistically significant. This is probably due to several features of our research design. First, even though we used 174 subjects in the experiments, the analysis focuses on only those subjects in the role of Player 1, reducing our sample size by one-half. The sample size is further reduced since we are interested in those subjects who had either a positive or negative experience in the first round, eliminating those who sent \$0 in the first round. While using deception would have allowed us to collect much more data, we felt that it was important to rely on a situation where subjects were actually matched with two other people in the room and this was transparent.

A second reason for not observing a larger effect may have to do with the subtlety of our Group treatment. In our experiment, trust in a collective entity is based solely on an individual's limited experience with members of a nominal group. By contrast, in an organizational context individuals typically have repeated experience with long-standing groups that strongly influence their members' lives. In such circumstances, it is reasonable to expect uniformity in behavior by group members. In our experiments, however, the 'group' was randomly determined by a median split of individuals' responses to an irrelevant and trivial guess. Therefore, it is striking that we observed *any* effect with such a slight group identity manipulation and we would expect an even larger effect in situations where the group or organizational identity is stronger.

The results are even more compelling when one considers that our group identity manipulation also likely created an 'out-group' bias, which would clearly work against our hypothesis. Individuals tend to view out-groups as less cooperative, honest, and trustworthy, and tend to expect less positive behavior from out-group members (Brewer 1979). Consequently, subjects interacting with counterparts categorized as members of an out-group would be biased toward viewing their counterparts as untrustworthy. Therefore, we expect that having subjects in the Group condition interact with two members of the same group – without it being a counterpart to their own group – might produce more striking results.

These limitations notwithstanding, we believe that this study makes a number of important contributions to research on trust in the economics, strategy and organizational literatures. Most importantly, the results of our experiment reinforce, and validate in a more carefully controlled setting, the finding by Zaheer et al. (1998) that trust in individuals and trust in collective entities are related but distinct. This suggests not only that it is meaningful to conceptualize the placing of trust in a collective entity, but also that

such trust may influence economic activity over and above individual trust. Consequently, it is important to carefully consider which level of analysis is most relevant when theorizing about the role of trust in the organization and coordination of economic activity. Further, recognizing that trust in a collective entity has a basis in group identification (Kramer et al. 1996) is essential.

We also go beyond earlier empirical research on trust in collective entities by highlighting trust transfer as an underlying causal mechanism that links trust in individuals and trust in collective entities. The evidence reported here is consistent with the idea that individuals use membership in a collectivity as a heuristic for determining the trustworthiness of members with whom they have no prior knowledge of or experience. This finding is striking because it suggests that the effects of an initial experience with a given representative of a collectivity extend beyond that relationship to interactions with other members of the collectivity. New relationships and interactions with previously unknown members of a collectivity do not start from a clean slate, but are construed through the lens of shared group identity with those with whom one has prior experience.

This study also makes a valuable empirical contribution by extending the trust game research paradigm. By embedding exchange relationships within the context of collective entities we are able to broaden the application of the trust game to a wider and more diverse set of phenomena that are more closely related to actual economic organizations and activities (cf. Buchan et al. 2002). Future research can draw on the research design developed here to address other questions involving trust in collective entities.

While this study advances our understanding of the relationship between trust in individuals and trust in collective entities, it also raises a number of important questions for future research. In particular, understanding the conditions that accelerate, alter or prevent the process through which trust transfers between individuals and collectivities represents a fruitful area of inquiry. For instance, in an organizational context, do certain structures, processes or incentives fundamentally alter the degree to which individuals rely on group identity as a heuristic for formulating initial trust impressions? A related and equally important question would be ascertaining the conditions under which group identity represents a useful and efficient heuristic versus an erroneous and costly bias. To the extent that these different circumstances can be identified, we would also want to gain insight into whether group identity as a basis for trust can be actively managed, produced or discouraged.

In sum, this research supports the idea that trust in a collective entity is related to, but distinct from, trust in the individual members of that collectivity. The findings are consistent with the view that economic actors develop perceptions about the trustworthiness of collective entities based on exchanges conducted with individual members of the collectivity. This trust in the collective entity then serves as a heuristic for trust in individuals where prior history or knowledge of members of the collectivity is limited or absent.

Acknowledgements

We would like to thank Robin Dawes and participants in the Organizational Behavior and Theory seminar series at Carnegie Mellon University and at the 2003 Economic Science Association meetings for useful comments on earlier versions of this chapter.

Notes

1. For similar arguments in sociology see Granovetter (1985) and Macauley (1963).
2. While we acknowledge differences in the object of trust, we do not consider differences in the origin of trust. Specifically, collective entities placing trust in (e.g. groups or organizations trusting) individuals or other collective entities is beyond the scope of this chapter (for a thoughtful discussion of this issue see research by Currall and Inkpen Chapter 13 this volume; 2002).
3. However, the average returns for sending \$5 and \$10 were \$7.17 and \$10.20, respectively. Berg et al. argue that the higher returns for these two amounts may reflect social norms concerning behavior towards players who sent half or the entire possible amount.
4. A plausible interpretation of this 'group effect' is that trusting behavior is normative, in the sense that it is part of a script that is primed by the experimental situation. If trusting behavior is primed, it will be rather insensitive to the object of trust, be it a specific person of a group member.
5. In the event that Player 1 had sent \$0, Player 2 did not need to make a choice, but we still required them to circle 'no choice' on the sheet so it would not be apparent who had received \$0 from their failure to circle something on the sheet.
6. Since we recruited from two large universities (total student populations 8000 and 16000) and from a large list of potential subjects, most participants did not know each other. Among the few that appeared to recognize someone in the room, it was very unlikely that they would be matched with the person they knew.
7. The gender composition did not differ significantly between the Control (30 percent female) and Group (39 percent female) conditions. Moreover, there are no gender differences in how much the first player sent to the second player in Round 1. Therefore, we omit further analysis of gender.
8. The change is brought about by an increase in the Group condition of the proportion of those sending money that send \$4. In the Control condition the proportion of those sending money that send \$4 remains roughly the same.
9. We can also use Goodman's (1964) test of 2 x 2 x 2 contingency tables to look for our hypothesized relationship between amount sent in Round 2, history and condition. Looking only at decisions to send \$4 versus a smaller amount in the second round (which is natural given the near 50-50 split of first-round choices using such categories), we find that we can reject the null hypothesis at $p < 0.1$. This is even though we change one value of 0 to 1 in order to perform this test, making the test more conservative.
10. Our results suggest an apparent asymmetry in that they are stronger when trust is abused than when it is honored, indicating that individuals may react to collective entities more strongly when a member violates their trust. However, there are at least two other possible explanations for such an asymmetry. First, almost everyone who trusts in the first round and has their trust honored, trusts again in the second round. Therefore, while there may be a difference in the propensity to trust in Round 2 between subjects in the Group and Control conditions, we might not observe it in these data due to such a 'ceiling effect'. Second, work on loss aversion (e.g. Kahneman and Tversky, 1979) indicates that outcomes that fall below a reference level carry more weight than those that exceed it. Following from this, one might expect that having trust abused (and ending up with a 'loss') might have a stronger effect than having it honored (and ending up with a 'gain').
11. The results are substantively unchanged if we use ordinary least squares instead of ordered logit.
12. All 14 subjects in the Group treatment who had trust honored in Round 1 sent \$4 in Round 2 (see Table 3.4). Therefore, behavior in this cell is perfectly identified in the second regression in Table 3.5. Consequently, we changed one such observation to \$2 in order to conduct the estimations for the fourth and fifth models in Table 3.5. This change works against our hypothesis, making the results more conservative. The results for the fourth model are exactly the same for any of the 14 values we could change. In column 5, we report the results of the regression using the change that produced the best fit (highest log-likelihood). However, the substantive results are unchanged for any of the 14 possible changes (i.e. the coefficients for the first and third independent variables are always significant at $p < 0.1$, while the coefficients for the other three independent variables are never statistically significant).

References

- Arrow, K. (1974), *The Limits of Organization*, New York: Norton.
- Ashraf, N., I. Bohnet and N. Piankov (2005), 'Measuring trust and trustworthiness', working paper.
- Bachmann, R. (2001), 'Trust, power and control in trans-organizational relations', *Organization Studies*, **22**, 337-65.
- Bachmann, R. (2003), 'The coordination of relations across organizational boundaries', *International Studies of Management & Organization*, **33**, 7-21.
- Barney, J.B. and M.H. Hansen (1995), 'Trustworthiness as a source of competitive advantage', *Strategic Management Journal*, **15** (special issue), 175-90.
- Berg, J., J. Dickhaut and K. McCabe (1995), 'Trust, reciprocity, and social history', *Games and Economic Behavior*, **10**, 122-42.

- Bradach, J.L. and R.G. Eccles (1989), 'Price, authority, and trust: From ideal types to plural forms', *Annual Review of Sociology*, **15**, 97–118.
- Brewer, M.B. (1979), 'In-group bias in the minimal intergroup situation: A cognitive-motivational analysis', *Psychological Bulletin*, **86**, 307–24.
- Bromley, P. and L.L. Cummings (1995), 'Transaction costs in organizations with trust', in R. Bies, B. Sheppard and R. Lewicki (eds), *Research on Negotiation in Organizations*, Greenwich, CT: JAI Press.
- Buchan, N., R. Croson and R. Dawes (2002), 'Swift neighbors and persistent strangers: A cross-cultural investigation of trust and reciprocity in social exchange', *American Journal of Sociology*, **108**(1), 168–206.
- Buchan, N., R. Croson and E. Johnson (in press), 'Trust and reciprocity: An international experiment', *Journal of Economic Behavior and Organization*.
- Corazzini, R. (1977), 'Trust as a complex multi-dimensional construct', *Psychological Reports*, **40**, 75–80.
- Currall, S.C. and A.C. Inkpen (2002), 'A multilevel approach to trust in joint ventures', *Journal of International Business Research*, **33**(3), 479–95.
- Giddens, A. (1990), *The Consequences of Modernity*, Stanford, CA: Stanford University Press.
- Goodman, Leo A. (1964), 'Simple methods for analyzing three-factor interaction in contingency tables', *Journal of the American Statistical Association*, **59**, p. 405.
- Granovetter, M. (1985), 'Economic action and social structure: The problem of embeddedness', *American Journal of Sociology*, **91**(3), 481–510.
- Güth, W., P. Ockenfels and M. Wendel (1997), 'Cooperation based on trust: An experimental investigation', *Journal of Economic Psychology*, **18**, 15–43.
- John, G. (1984), 'An empirical investigation of some antecedents of opportunism in a marketing channel', *Journal of Marketing Research*, **21**, 278–89.
- Kahneman, D. and A. Tversky (1979), 'Prospect theory: An analysis of decisions under risk', *Econometrica*, **47**, 313–27.
- Kramer, R.M., M.B. Brewer and B.A. Hanna (1996), 'Collective trust and collective action: The decision to trust as a social decision', in R.M. Kramer and T.R. Tyler (eds), *Trust in Organizations: Frontiers of Theory and Research*, Thousand Oaks, CA: Sage.
- Lewis, J.D. and A.J. Weigert (1985), 'Trust as social reality', *Social Forces*, **63**(4), 967–85.
- Luhmann, N. (1979), *Trust and Power*, Chichester: Wiley.
- Macauley, S. (1963), 'Non-contractual relations in business', *American Sociological Review*, **28**, 55–67.
- Mayer, R.C., J.H. Davis and F.D. Schoorman (1995), 'An integrative model of organizational trust', *Academy of Management Review*, **20**, 709–34.
- McEvily, B., V. Perrone and A. Zaheer (2003), 'Trust as an organizing principle', *Organization Science*, **14**, 91–103.
- Powell, W.W. (1990), 'Neither market nor hierarchy: Network forms of organization', in L.L. Cummings and B. Staw (eds), *Research in Organizational Behavior*, Greenwich, CT: JAI Press, vol. **12**, 295–336.
- Rotter, J.B. (1967), 'A new scale for the measurement of interpersonal trust', *Journal of Personality*, **35**, 651–65.
- Rousseau, D.M., S.B. Sitkin, R.S. Burt and C. Camerer (1998), 'Not so different after all: A cross-discipline view of trust', *Academy of Management Review*, **23**, 292–404.
- Snijders, C. and G. Keren (1998), 'Determinants of trust', in D.V. Budescu, I. Erev and R. Zwick (eds), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, Mahwah, NJ: Lawrence Erlbaum.
- Stewart, K.J. (2003), 'Trust transfer on the world wide web', *Organization Science*, **14**, 5–17.
- Tajfel, H., M. Billig, R.P. Bundy and C. Flament (1971), 'Social categorization and intergroup behavior', *Journal of Experimental Social Psychology*, **1**, 149–77.
- Van Huyck, J.B., R.C. Battalio and M.F. Walters (1995), 'Commitment versus discretion in the peasant-dictator game', *Games and Economic Behavior*, **10**, 143–70.
- Williamson, O.E. (1985), *The Economic Institutions of Capitalism*, New York: The Free Press.
- Zaheer, A., B. McEvily and V. Perrone (1998), 'Does trust matter? Exploring the effects of interorganizational and interpersonal trust on performance', *Organization Science*, **9**(2), 141–59.
- Zucker, L.G. (1986), 'Production of trust: Institutional sources of economic structure', in L.L. Cummings and B.M. Staw (eds), *Research in Organizational Behavior*, Greenwich, CT: JAI Press, vol. **8**, 53–111.