

UNDERSTANDING TRUST*

Paola Sapienza, Anna Toldra-Simats and Luigi Zingales

The World Values Survey (WVS) question on trust has been widely used to study the economic effect of trust. Recent work, however, questions its validity as an accurate measure of trust by showing that it is not correlated with the sender's behaviour in the Berg *et al.* trust game. What measure then should we trust to measure trust? In this article, we argue that the sender's behaviour in a trust game is driven both by beliefs and by preferences. In contrast, WVS-like measures capture mostly the belief-based component of a trust game.

In the last 15 years, economists have increasingly paid attention to the role trust plays in economic activity. From economic growth (Knack and Keefer, 1997) to size of firms (La Porta *et al.*, 1997; Bloom *et al.*, 2009), from financial development (Guiso *et al.*, 2004, 2008) to international trade and investments (Guiso *et al.*, 2009), many economic phenomena have been related to the level of trust. But what is trust? And how do we measure it?

'When we say we trust someone or that someone is trustworthy – writes Gambetta (2000) – we implicitly mean that the probability that he will perform an action that is beneficial (...) is high enough for us to consider in engaging in some form of cooperation with him.' In Gambetta's (2000) definition trust is a belief, which can be measured as a probability. Given this interpretation, it would seem natural to measure it using the Berg *et al.* (1995) game, also known as the 'trust' game.

As a measure of trust, however, most papers in this literature use the answer to the World Values Survey (WVS)/General Social Survey (GSS) question 'Generally speaking, would you say that most people can be trusted or that you cannot be too careful in dealing with people?' Despite its broad use, it is not entirely clear what this question exactly measures, nor what is the correlation between the WVS-question and experimental measures of trust. Both Glaeser *et al.* (2000) and Lazzarini *et al.* (2003) have shown that the answers to the WVS-question are not correlated with the sender's behaviour in the standard trust game, but they are correlated with the receiver's behaviour in the same game. These papers conclude that the WVS-question is a measure of trustworthiness and not of trust. On the other hand, Fehr *et al.* (2003) and Bellemare and Kroeger (2007) challenge this result. Using a large sample of German and Dutch households, respectively, these studies show that the sender's behaviour in a

* Corresponding author: Anna Toldra-Simats, Universidad Carlos III de Madrid. Calle Madrid 126, Edificio 7. Getafe 28903, Madrid, Spain. Email: atoldra@emp.uc3m.es

This study is part of a larger project funded by the Templeton Foundation. Without its support none of this would have been possible. In addition, Paola Sapienza thanks the Zell Center, and Luigi Zingales thanks the Center for Research in Security Prices (CRSP) and the Initiative on Global Markets at the University of Chicago Booth School of Business for financial support. Anna Toldra-Simats thanks the Kellogg School of Management (Northwestern University) because this is where this study started. We thank Ernesto Reuben for valuable suggestions throughout the process of writing this article, and Janice Luce and Peggy Eppink for editorial help. We also thank three anonymous referees.

trust game is correlated with other survey-based measures of trust, which in turn are not correlated with trustworthiness.

These contradictory findings raise the question of whether survey-based measures of trust are good indicators. More specifically, given its pre-eminence in the literature, what is the WVS-question measuring? Is it trust, trustworthiness or neither of the two? Can we dismiss the WVS-question based on the assumption that the Berg *et al.* (1995) trust game and its variations provide an accurate measure of trust? Or should we reject the trust game as a measure of trust since, as Levitt and List (2007) claim, in the laboratory there are several factors, such as scrutiny by the experimenter, that distort subjects' behaviour? In other words, can we trust the trust game and/or the WVS-question to accurately measure trust?

Our starting point is that neither method is a perfect indicator of trust as defined by Gambetta (2000). To begin with, the sender's behaviour in the trust game is affected by other motivations besides the sender's belief in the receiver's trustworthiness, such as individual risk aversion (Karlan, 2005; Schechter, 2007) and other-regarding preferences like altruism (Cox, 2004; Ashraf *et al.*, 2006). In the light of these results, we understand the act of trusting as the combination of two components: beliefs in other people's trustworthiness and the specific preferences of the sender (risk aversion, inequality aversion, altruism).¹ It is useful to keep these two components of trust separate because their persistence across different situations might be different; their persistence outside the laboratory might be different (as Levitt and List, 2007 argue, in the laboratory we observe a lot of pro-social behaviour because subjects feel they are under the experimenter's scrutiny and thus behave in a more socially acceptable way). Since the ultimate goal is to find a measure of trust that captures some general attitudes of a population, it is important to find out what component(s) of the trust game (if any) measure trust so defined and what the WVS-trust question measures, as well as the relationship between the two.

To this end, in this article, we run a modified trust game in which all participants play both as senders and as receivers. As in the traditional trust game, senders are initially endowed with \$50 and are asked to send an amount between \$0 and \$50 to the receiver. The amount sent is then tripled by the experimenter and the receiver is asked to decide how much of it to return to the sender. In addition to this standard component, we ask the sender to report his beliefs about the receiver's behaviour for every possible amount sent, using the strategy method. By doing so, we can separate the sender's expectations about the behaviour of the receiver (the beliefs component) from his actions, which are affected also by his utility function. In addition, the use of the strategy method allows us to compare the trust behaviour of senders for different amounts sent.

We find that the quantity sent in the trust game is not only correlated with the sender's expectation of the receiver's trustworthiness but also with his preferences. Interestingly, we find that beliefs are significantly correlated with the quantity sent only when subjects send more than 25% of their initial endowment. This is interesting because in a trust game like ours where only senders are initially endowed and the

¹ For a distinction between preferences and beliefs in the related literature, see Bohnet and Zeckhauser (2004) and Bohnet *et al.* (2008).

amount sent is tripled, sending an amount to the responder lower or equal to 25% of the initial endowment can be interpreted as an act of charity (i.e. not involving beliefs in others' trustworthiness) rather than an act of trust.

We then turn to examine the WVS-trust question. We find that senders' expectations are correlated with the WVS-trust question, as well as other attitudinal questions on trust. This suggests that the WVS-question captures the expectation component of the trust game. Most interestingly, expected trustworthiness is only correlated with the WVS-question when the sender is calculating the expected amount returned for large amounts sent. This result further supports the idea that for smaller amounts of money sent, the sender's decision does not involve trust.

Then, we analyse the correlation between the sender's expectations of the receiver's trustworthiness with the sender's actual trustworthiness when he plays as a receiver. We find that players extrapolate their opponent's behaviour from their own. This finding can explain why in some cases the WVS-question is correlated with trustworthiness.

In summary, neither the WVS-question nor the trust game measures trust as defined by Gambetta properly. The best measure one can obtain from the trust game is not the amount sent, but the expectation about the amount returned for large amounts sent, since this variable is the least contaminated by other considerations.

The rest of the article proceeds as follows. Section 1 describes the experimental design, the subject pool and the survey and presents summary statistics. In Section 2, we analyse senders' behaviour in the trust game and how it relates to their expectations and preferences. In Section 3, we analyse the relationship between the WVS-trust question and the trust game. Section 4 presents a discussion of related issues, and Section 5 concludes.

1. Experimental Design and Survey

The data for this experiment were collected as part of the Chicago-Templeton MBA Longitudinal Study (CTMLS). All the full-time MBA students of the 2008 class at the University of Chicago's Booth School of Business were asked to complete a survey and take part in a laboratory experiment as part of a mandatory class. While participation was mandatory, the Institutional Review Board at the University of Chicago required that the subjects be offered the opportunity to opt out from the study by not consenting to the use of their data for research purposes. Of the 550 MBA students, 548 completed the survey and 544 played the games; 502 (92.28%) of the 544 consented to the use of both their survey and experimental data. For this article, we use data from these two sources in addition to admission data obtained from the school (also with the students' consent). Each of these data sets is described below.

1.1. *Survey and Subject Pool*

Completing the survey was compulsory and was done online. It was sent to the students 2 weeks before the games took place and the deadline was the start of the experiment. We kept track of the dates and times at which students completed the survey. On average, students completed the survey 7.33 days before participating in the games.

The survey was designed to acquire demographic data and measure various personality traits (all the questions asked in the survey are available in Reuben *et al.*, 2008). In this article, we use several variables from the survey that aim to measure trust. Our main survey measure of trust is the standard WVS question: ‘Generally speaking, would you say that most people can be trusted or that you cannot be too careful in dealing with people?’ We provided three answers for subjects to choose from: most people can be trusted; Cannot be too careful in dealing with people; Do not know. Table A1 in the online Technical Appendix shows summary statistics about this variable: a bit more than half of the subjects in our sample (53.59%) answered that most people can be trusted.

The survey included two additional questions regarding trust attitudes that are useful for our analysis as additional measures of trust. The first: ‘Suppose that a new and very desirable dorm/apartment has become available. The University of Chicago organises a lottery to assign it among the many applicants. How confident are you that the allocation will be fair?’ The choice of answers was: not at all; not much; quite a lot; a great deal; I do not know. Almost 90% of the students answered that they trust the University of Chicago quite a lot (41.04%) or a great deal (45.62%). Only 0.79% do not trust it at all. The second question was ‘Suppose that while walking on Michigan Avenue in Chicago you lose your wallet with \$1,000 dollars inside. A random person that you do not know finds it. He or she does not know you, but he or she is aware that the money belongs to you and knows your name and address. He or she can keep the money without incurring any punishment. According to you, what do you think is the probability he or she will return the money to you? (Report a number between 0 and 100, where 0 means that the money would not be returned for sure and 100 means that it will be returned for sure.)’. On average, students thought they had a 34.82% chance of getting their wallet back. The mode is at 50% with 95 students; 37% of the students thought that the probability was less than 25%. Additional summary statistics for these variables can be found in Table A1 in the online Technical Appendix.

Table A2 in the online Technical Appendix provides pair-wise correlations between the three trust variables. Both the wallet question and the Chicago question are positively and statistically significantly correlated with the WVS-trust question and with each other.

We also use information provided by the admissions office of the University. It includes additional demographic characteristics of the subjects such as their age, gender and GMAT score. We will use these variables as controls in our analysis. Table A3 in the online Technical Appendix shows summary statistics of these variables. The average age in our sample is 28.3 years, 31% of subjects are female, and their average GMAT score is 705.

1.2. *Laboratory Experiments*

The laboratory experiments consisted of two lotteries, four games and an auction played in the following order: lottery with losses, asset market game, trust game, competition game, chocolate auction, social dilemma game and lottery without losses. The experiment was programmed in z-Tree (Fischbacher, 2007) and played in four sessions in four large classrooms. We paid students by randomly drawing one game/

lottery at the end of the experiment. In total, 544 MBA students participated in the experiment and earned on average \$78.32 in addition to a \$20 show-up fee. The experiment lasted approximately one and a half hours.

In this article, we concentrate on the trust game, the social dilemma game and the lottery game without losses. The trust game is the main object of this study. We use the social dilemma game and the lottery game as explanatory variables in our analysis. Below is a description of these games. For a description of the other games, see Reuben *et al.* (2008).

1.2.1. *The trust game*

The trust game we used is a slightly modified version of the trust game initially designed by Berg *et al.* (1995), henceforth BDM, but also widely used in the experimental literature. In this game a sender is endowed with an amount of money y . The sender decides how much to send, $s \in [0, y]$, to a receiver. Any amount sent is multiplied by 3. The receiver then decides how much to return, $r \in [0, 3s]$, to the sender. Consequently, the pay-off of the sender equals $y - s + r$, and that of the receiver equals $3s - r$. The amount sent is frequently referred to as a measure of trust, and the amount returned as a measure of trustworthiness.

In our experiment, subjects were asked to make three decisions: first they all played the trust game in the role of the sender, then we elicited their beliefs about the behaviour of the receiver, and third they played the trust game in the role of the receiver. The screenshots that students faced when playing the trust game can be found in Appendix B in the online Technical Appendix. Subjects' earnings were determined by randomly selecting one of the three decisions at the end of the experiment. When making a decision, subjects did not know what future decisions they would be asked to make; however, subjects did know that they would make three decisions and it was emphasised that the three decisions were independent in the sense that their choices in one decision would not affect their earnings in subsequent decisions. In between decisions, no feedback was given with respect to the behaviour of other subjects.²

In the first decision, subjects played as senders and decided how much of their initial endowment of \$50 they wanted to send to the receiver. They could send any amount between \$0 and \$50 in multiples of five. In the second decision, subjects indicated how much they expected the receiver to return for each possible amount sent (filling an array like the one in screenshot 2 in Appendix B in the online Technical Appendix). To motivate subjects to answer accurately, their earnings in this decision consisted of \$10 for each amount sent when the distance between their expectation and the actual response was within 10% of the amount received (i.e. if $r - 0.1 \times 3s \leq E[r] \leq r + 0.1 \times 3s$).³ In the third decision, subjects played as receivers and were

² Due to the logistics of running such a large experiment, the three decisions were made sequentially in the same order by all subjects. Note that since no feedback was given, subjects made their decisions with the same information.

³ Incentivised procedures to elicit the mean of a distribution require complicated elicitation techniques due to potential differences in risk aversion and probability weighting (Manski, 2002; Offerman *et al.*, 2009). For this reason, we opted for a cognitively simpler elicitation procedure that would not distract subjects from the game they were playing.

asked to indicate the amount they were willing to return for every possible amount sent, that is, they made their choices using the strategy method (Selten, 1967).⁴ To facilitate calculations during their third decision, subjects could use two buttons to calculate how their decision would affect their earnings and the earnings of their opponent.⁵

Hence, each subject played the trust game twice, once in the role of the sender and once in the role of the receiver. Subjects were randomly re-matched so when they played as senders and receivers they played with a different person. Subjects were not informed of the identity of the people they played with, which when combined with the large sample size ensures that their decisions were anonymous.

Subjects in our sample sent on average \$18.82, which is 37.64% of senders' initial endowment. The standard deviation of the amount sent was \$14.9. Figure C1 in the online Technical Appendix displays the distribution of quantities sent. Table C1 in the online Technical Appendix reports additional summary statistics of subjects' behaviour in the trust game. The second and third columns of this Table show, respectively, the average return and the average return proportional to the amount received for every quantity sent (displayed in the first column). The fourth and fifth columns display, respectively, average and proportional expected returns also conditional on the quantities sent. Both the returns and the returns proportional to the amount received are strictly increasing with the amount sent. That the expected return as a proportion of the quantity sent is systematically 4–6 percentage points above the actual return as a proportion of the amount received suggests that subjects were over optimistic about the receiver's trustworthiness.

1.2.2. *The lottery without losses*

We use the lottery without losses to measure subjects' risk aversion in small gambles. In a similar way as Holt and Laury (2002), we asked subjects to choose 15 times between two options: Option *A* and Option *B*. When choosing Option *A*, subjects could win an amount of dollars with certainty, ranging from \$50 up to \$120 in increments of five. Option *B* always consisted of a lottery offering \$200 or \$0 with equal probability. At the end of the game, one of the 15 choices was randomly chosen and subjects were paid according to their decision.⁶

According to the pay-offs in this game, an extremely risk-averse individual should choose Option *A* in all settings, whereas an extreme risk seeking individual should always choose Option *B*. For those in-between, as the certain amount increases, a subject should cross over from Option *B* to Option *A*. The less risk-averse the subject is, the later the switch will occur. A risk-neutral individual should choose Option *B* until

⁴ The strategy method has been criticised for it may elicit strategies that differ from those used in a strictly sequential environment. However, in a recent survey, Brandts and Charness (2011) analyse around 30 papers and find that evidence mostly suggests that the strategy method and the direct-response method produce similar results. If there are any differences, one can observe a difference in games that involve punishment, but not in sequential social dilemmas like the trust game. Also see Brandts and Charness (2000) and Vyrastekova and Onderstal (2005) for additional discussion on the strategy method.

⁵ The three screenshots of the trust game that the students were presented with can be found in Appendix B in the online Technical Appendix.

⁶ The screenshots of this game can be found in Reuben *et al.* (2008).

Option *A* is worth \$95 or \$100 and then switch to Option *A*. We refer to the last value of Option *A* for which the subject chooses Option *B* as the subject's certainty equivalent,⁷ which is a measure of the subject's risk aversion provided that the subject changed options at most once.⁸

Figure C2 in the online Technical Appendix displays the number of risky choices made by subjects in our sample. On average, subjects made seven risky choices (standard deviation is 3.5) corresponding to an average certainty equivalent of \$80. The mode is 10 risky choices (certainty equivalent \$95) with 107 students (21.31% of the sample). Only 21 students (2.93% of the sample) made more than 11 risky choices exhibiting risk loving behaviour.

1.2.3. *The social dilemma game*

Subjects played a social dilemma game based on the commonly used linear public good game (Marwell and Ames, 1981; Isaac *et al.*, 1984). Subjects were randomly assigned to groups of eight and given an endowment of \$50. Each subject then decided whether to contribute c to the public good. Contributions to the public good are costly to the subject but increase the earnings of others. Specifically, subject i 's earnings equal $\$50 - c_i + 0.3 \times \sum_j c_j$. Unlike in most public good experiments, the contribution decision was binary: subjects could contribute either all their endowment or nothing, $c \in \{\$0, \$50\}$. Note that overall pay-offs are maximised if all eight subjects contribute \$50. However, since an individual receives only \$15 for her \$50 contribution, she maximises her monetary pay-off by not contributing.

The experiment was designed to elicit the willingness of subjects to conditionally cooperate. For this purpose, we used a variation of the design employed by Fischbacher *et al.* (2001). Subjects made two contribution decisions: first an 'unconditional' decision and after that a 'conditional' one. The unconditional decision was simply to either contribute the \$50 to the public good or not. For their conditional decision, we used the strategy method (Selten, 1967) to allow subjects to condition their contribution on the number of group members contributing to the public good. Specifically, subjects had to indicate whether they would contribute their \$50 if x other group members also contributed theirs, and x ranged from 0 to 7.⁹ To determine each subject's pay-off, one of the two decisions was randomly selected. If the unconditional decision was chosen, then that subject's pay-off was given by her unconditional decision, the unconditional decision of the six other group members and the conditional decision of one group member. If the conditional decision was chosen, the subject's pay-off was given by her conditional decision and the others' unconditional

⁷ In reality, the certainty equivalent would fall within \$5 of this number. That is, if a subject switches to the safe choice at \$100, the certainty equivalent falls between \$95 and \$100.

⁸ The risk aversion measure was set to missing for 24 subjects because they switched options more than once. In addition, we assigned a certainty equivalent of \$45 to the 42 subjects who always chose the safe option and one of \$120 to the nine subjects who always chose the risky option. For robustness, we run several tests assigning different values to the subjects in the tails. For example, in another specification we set the certainty equivalent to \$30 for those 42 subjects who always chose the safe option and one of \$150 for the nine subjects who always chose the risky option; in a different specification we set the certainty equivalent to \$20 and \$200 respectively. With more extreme values, the coefficient of risk tolerance decreases slightly, but the results of our regressions do not change qualitatively.

⁹ The screenshots of this game can be found in Reuben *et al.* (2008).

decision. All subjects made both decisions without knowing what others in their group did. Furthermore, when making their unconditional decision, subjects were not aware their second decision would be the conditional one.

We use subjects' conditional cooperation choices in the social dilemma game to measure other-regarding preferences.¹⁰ This measure includes several motivations of social preferences since, as stated by Fischbacher *et al.* (2001), conditional cooperation can be considered as a consequence of some fairness preferences like 'altruism', 'warm-glow', 'inequity aversion' or 'reciprocity'.¹¹ Table C2 in the online Technical Appendix shows summary statistics of subjects' cooperation conditional on other participants' cooperation. On average, subjects tend to cooperate more the more other subjects cooperate in their group. However, when inspecting the data at the individual level, we observe that subjects are heterogeneous. If we use the classification of Fischbacher *et al.* (2001), we find that 232 subjects (46.59%) exhibit 'free riding' behaviour, that is, they never cooperate; 246 subjects (49.40%) are (weakly) monotonic conditional cooperators, that is, they start to cooperate as more subjects cooperate, with 24 of them (i.e. 4.82% of the sample) cooperating only when the seven other subjects in the group cooperate; and 20 of them (4%) exhibit hump-shaped contributions meaning that they start to cooperate as more subjects cooperate in their group but they stop cooperating after a higher number of other subjects cooperate.

For our analysis in subsequent Sections, we use the subjects' conditional strategy as a measure of their other-regarding preferences. Specifically, we have created a dummy variable that equals the number of times a subject cooperates in the conditional cooperation setting. For the sake of brevity, we report regressions using only this definition of other-regarding preferences. However, our results are not sensitive to defining this variable differently.¹²

2. Sender's Behaviour in the Trust Game

2.1. What Does the Trust Game Measure?

If we want to use the trust game to 'validate' the WVS-question, we need to be clear about which aspect of trust we want to measure. According to Gambetta (2000), when we trust an individual 'we implicitly mean that the probability that he will perform an

¹⁰ Since our objective is to distinguish preferences and beliefs in the subjects' behaviour, we disregard the unconditional cooperation choice because this decision may be based, in addition to preferences, on beliefs about other subjects' willingness to cooperate.

¹¹ As in the basic models of social preferences, that is, dictator game, cooperation game, our setting is not designed to distinguish among different types of other-regarding preferences. In addition to that, efficiency gains, that is, the fact that the amount sent is tripled (or doubled), may also motivate the sender to send positive amounts if she trusts that the receiver is willing to share such gains (Charness and Rabin, 2002; Ashraf *et al.*, 2006).

¹² We replicate the results of the article using the following four alternative measures of other-regarding preferences: a dummy variable that equals 0 if a subject never cooperates and equals 1 otherwise; a dummy variable equal to the previous one but excluding the 20 subjects who show hump-shaped behaviour; another dummy that takes the value of 0 if a subject never cooperates or if she cooperates only if seven other subjects cooperate, and 1 otherwise; and a dummy that equals the last one but excludes the 20 subjects who show hump-shaped behaviour. The additional regressions are available upon request from the authors.

action that is beneficial...is high enough for us to consider in engaging in some form of cooperation with him'. Similarly, James (2002) defines trust as an expectation. 'To say "A trusts B" means that A expects B will not exploit a vulnerability A has created for himself by taking the action'.

Both these definitions focus on an individual's expectation of another individual's trustworthiness. In Gambetta's case, this should be sufficiently high to induce an individual to cooperate, for James it should be high enough so that the subject feels not taken advantage. As such, both definitions seem in line with what the WVS question tries to measure. The WVS-trust question aims to measure generalised trust; that is, the beliefs of the respondent regarding the trustworthiness of others in general.

At first sight, these definitions seem also in line with the behaviour of the sender in the Berg *et al.* (1995) trust game. After all, why should a subject send part or all of his endowment if he does not trust a receiver to cooperate (i.e. not behave in an opportunistic way)?

Yet, recent research has pointed out that the sender's behaviour in the trust game is affected by several other factors in addition to trust, starting with risk aversion (Karlan, 2005; Schechter, 2007). Even if individuals should be risk neutral for small gambles, they are often not (Holt and Laury, 2002). Hence, if we want to extract the trust component from the amount sent in a trust game, we should at the very minimum control for risk aversion.¹³

Furthermore, the senders' behaviour in the trust game is indicative of expectations in receivers' trustworthiness only if senders are selfish. However, if people were selfish they would always return zero when acting as receivers, which is neither the case in the trust games reported in the literature, nor in ours. In fact, it is remarkable that 90% of the MBA students in one of the most economically minded MBA programmes in the world actually returned a positive amount. Indeed, from its inception, researchers have interpreted the behaviour of the receiver in the trust game as an indication of other-regarding preferences.¹⁴ For example, with a triadic design, Cox (2004) finds evidence that receivers are motivated by both reciprocity and altruism.

If subjects have other-regarding preferences when they play as receivers, it is likely that they have other-regarding preferences when they play as senders. In our sample, when we tried to rationalise each sender's behaviour as an optimal choice given his expectation and his level of risk aversion, we could not make sense of it unless we hypothesised that the sender had some form of other-regarding preferences.¹⁵ We thus expect senders in our trust game to exhibit other-regarding preferences.

Other-regarding preferences may in fact be particularly important in the trust game studied in this article. As explained before, we use a modified version of the typical BDM-trust game. The major difference between the BDM game and the one that we

¹³ In a related paper, Houser *et al.* (2010) find that risk attitudes do not predict decisions in the trust game.

¹⁴ The following are some of the models that can explain why receivers return money: Rabin (1993), Levine (1998), Fehr and Schmidt (1999), Bolton and Ockenfels (2000), Charness and Rabin (2002), Falk and Fischbacher (2006) and Cox *et al.* (2008).

¹⁵ For example, individual choices in our data can be explained as rational ones if we assume – following Fehr and Schmidt (1999) – that a sender's expected utility depends not only on his pay-off but also on the comparison between his pay-off and his opponent's pay-off. This analysis is available from the authors.

and many others use is that in this version only the sender is initially endowed with money, whereas in the BDM version both the sender and receiver are initially endowed. When receivers are not endowed, egalitarian senders will send money to receivers regardless of whether they trust them or not, that is, they act based only on their other-regarding preferences. Therefore, we need to confirm the extent to which senders in our trust game send money to receivers as an act of trust. Finally, by the definition of trust, we expect senders' elicited expectations about others' returns to also determine the quantity sent.

Finally, there is another difference between the traditional trust game, the two above-mentioned definitions and the WVS-question. In the traditional trust game, subjects play anonymously. In contrast, there is nothing in the WVS-question nor in the above definitions that requires anonymity. In fact, the typical real-world situations people might think of when answering the WVS-question might be non-anonymous situations.¹⁶ This is another likely reason why the answers to the WVS-question might not be highly correlated with experimental measures of trust derived from the Berg *et al.* (1995) trust game.

2.2. Empirical Analysis of the Sender's Behaviour

The sender's behaviour in the trust game is driven by two factors: preferences and beliefs. For a given type of preferences, a sender who has higher expectations about other people's trustworthiness will send more. Similarly, for a given level of expectations, a more altruistic sender (or a less risk-averse sender) will send more.

On the basis of the above definitions, we regard only the beliefs component as a measure of trust. To isolate it, we use the sender's expectations, which we elicited during the experiment. These expectations should be unaffected by the sender's utility function and should be a true measure of the sender's expected level of the receiver's trustworthiness – that is, trust.

Table 1 validates our hypothesis that the sender's behaviour is affected by these two factors. Our dependent variable is the quantity sent. Our explanatory variables are the subjects' preferences and her beliefs. As we elicited expectations using the strategy method, the amount expected back is conditional on the amount sent. The three panels of Table 1 use this information in different ways. Panel (a) includes only the expected return conditional on sending \$50. Panel (b) includes 10 different regressions with the expected returns conditional on each possible quantity sent. Panel (c) pools all the regressions of panel (b) introducing an individual fixed effect.¹⁷

¹⁶ We thank an anonymous referee for this important point.

¹⁷ The number of usable observations is 502. When including the risk tolerance measure, 24 observations are set to missing because these subjects switch options more than once. This brings us to 478 observations. In addition, 48 subjects answered 'do not know' to the WVS-trust question, this should bring us down to 430 subjects; however, four of the subjects who answered 'do not know' are also set to missing because they switch choices in the risk game more than once. Therefore, only 44 subjects should be subtracted when adding the WVS-trust question in the regression, which brings us down to 434. In panel (c), data are organised as a panel with 10 observations per subject corresponding to the possible quantities sent. The number of usable observations is 5,020 but 53 subjects (530 observations) are removed from the analysis because these subjects chose to send \$0 in the trust game and thus there is no variation in the dependent variable for these subjects, leaving us with 4,490 observations in the regression.

Table 1

Panel (a): *Quantity Sent to Second Mover*; Panel (b): *Quantity Sent to Second Mover – Conditional*; Panel (c): *Quantity Sent to Second Mover – Pooled Data*

	Panel (a)			Panel (b)			Panel (c)						
	(1)	(2)	(3)	if sent \$5	if sent \$10	if sent \$15	if sent \$20	if sent \$25	if sent \$30	if sent \$35	if sent \$40	if sent \$45	if sent \$50
Risk tolerance	0.18*** (0.04)	0.17*** (0.04)	0.17*** (0.04)	0.17*** (0.04)	0.17*** (0.04)	0.17*** (0.04)	0.16*** (0.04)	0.16*** (0.04)	0.16*** (0.04)	0.15*** (0.04)	0.16*** (0.04)	0.16*** (0.04)	0.16*** (0.04)
Other-regarding preferences	0.82*** (0.36)	0.73*** (0.33)	0.59* (0.32)	0.83*** (0.34)	0.83*** (0.34)	0.83*** (0.34)	0.83*** (0.34)	0.85*** (0.34)	0.83*** (0.34)	0.77*** (0.33)	0.78*** (0.33)	0.75*** (0.33)	0.74*** (0.32)
Trust-WVS	2.49* (1.42)	1.59 (1.36)		0.21*** (0.06)	0.24*** (0.05)	0.21*** (0.08)	0.21*** (0.06)	0.24*** (0.05)	0.21*** (0.04)	0.20*** (0.03)	0.17*** (0.03)	0.17*** (0.03)	0.15*** (0.02)
Expected trustworthiness of other (if sent \$50)		0.15*** (0.02)	0.15*** (0.02)	-0.39 (0.27)	-0.39 (0.27)	-0.39 (0.27)	-0.44* (0.26)	-0.48* (0.26)	-0.50* (0.26)	-0.53*** (0.26)	-0.55*** (0.26)	-0.55*** (0.26)	-0.56*** (0.26)
Constant	1.24 (3.32)	-5.44 (3.39)	-4.74 (3.23)	-1.59 (1.52)	-1.52 (1.52)	-1.55 (1.52)	-1.53 (1.49)	-1.47 (1.48)	-1.26 (1.48)	-1.29 (1.46)	-1.51 (1.46)	-1.36 (1.44)	-1.31 (1.44)
Observations	434	434	478	478	478	478	478	478	478	478	478	478	478
R ²	0.068	0.146	0.137	0.071	0.093	0.073	0.093	0.115	0.119	0.137	0.135	0.149	0.153
Expected returns													
Possible amount sent													
Constant													
Individual FE													
No. of observations													
R ²													

Note. Standard errors are in parentheses. *Significant at 10%. **Significant at 5%. ***Significant at 1%.

Column 1 of panel (a) reports an ordinary least squares regression of the dollar amount sent by subjects in the trust game on their risk tolerance (i.e. the certainty equivalent of the lottery), their other-regarding preferences (i.e. their willingness to conditionally cooperate in the social dilemma game) and their answer to the WVS-trust question.

As expected, more risk-tolerant individuals send more, as do subjects who are motivated by other-regarding preferences. Increasing the certainty equivalent by \$5 increases the amount sent by 90 cents on the dollar (i.e. 0.18×5); or alternatively, a one-standard deviation decrease in risk aversion leads to a 17% increase in the average amount sent.¹⁸ Also, people with other-regarding preferences send on average \$0.80 (4.25%) more to receivers.¹⁹ The WVS-trust variable has a positive and statistically significant effect at the 10% level. Subjects who respond that most people can be trusted send on average \$2.72 (13%) more to the receiver.

This effect drops almost in half and loses its statistical significance when we control for the senders' expected return (conditional on \$50 being sent). In contrast, the senders' expected return is highly economically and statistically significant (column 2). For each additional dollar the sender expects to receive back, she will increase the amount sent by 15 cents. As we will see below, the variable 'expected return' is correlated with the WVS-trust variable; this is the reason why its insertion eliminates the impact of the WVS-trust measure.

In Column 3, we run the same regression as before but omitting the WVS-trust variable. A one-standard deviation increase in the expected trustworthiness of others (conditional on \$50 sent) increases the average amount sent by 24%. According to these results, the quantity sent in the trust game is indeed a combination of two components: the preferences of the sender (risk aversion and other-regarding preferences) and his beliefs in others' trustworthiness (trust).²⁰ This second component seems to be better captured by the sender's expectations than by the WVS measure.

In Panel (b), we run a regression similar to the third column of panel (a), with the senders' expected returns conditional on all possible quantities sent. That is, the first column in panel (b) shows the results of the regression where the senders' expectations correspond to an amount sent equal to \$5, the second column shows the regression results where the senders' expectations correspond to an amount sent equal to \$10, and so on. We also include demographic characteristics as additional controls.

¹⁸ We also computed a constant relative risk aversion coefficient by assuming a utility function of the form $u(x) = x^{1-r}/(1-r)$, where r is the coefficient of relative risk aversion. We obtain an interval for the coefficient of risk aversion using the pay-offs of the different lotteries, the hypothesised utility function, and the crossover point from the risky to the riskless option of each subject. We use the mid-point of this interval as the relative risk aversion coefficient of each subject. When conducting the above analysis with the constant relative risk aversion coefficient, our results do not change.

¹⁹ Since our objective is to distinguish preferences and beliefs in the subjects' behaviour, we do not use the unconditional cooperation choice because this decision may contain, in addition to preferences, beliefs about other subjects' willingness to cooperate. Nonetheless, we have run a check with the unconditional cooperation variable and unconditional cooperation also seems to be correlated with the quantity sent in the trust game.

²⁰ In our model, we assume that beliefs are not correlated with social preferences. This is a standard assumption in the economic models of other-regarding preferences (Fehr and Schmidt, 2006).

Interestingly, the expected trustworthiness is significantly related to the quantity sent only for quantities sent of \$20 and above. In a trust game like ours, where only the sender receives an initial endowment and the amount sent is multiplied by 3, the sender gains more than the receiver as long as the sender sends \$12.50 or less. Thus, for these (low) amounts sent, sending money to the receiver can be interpreted as an act of charity more than an act of trust. This may be the reason why beliefs are not significantly correlated with the quantity sent for these amounts.

The coefficient and significance of the rest of variables stay the same in all these regressions, confirming that the behaviour of senders in the trust game combines beliefs and preferences. As for the controls, age has a negative and significant effect on the quantity sent, and subjects with higher GMAT scores send significantly more. Females tend to send less, but the coefficient is not significant.

Panel (c) displays the results of a regression model where the data are organised as a panel with 10 observations per subject corresponding to the 10 possible quantities sent: (\$5, \$10, \$15, ..., \$50). In this framework, behaviour is predicted by combining the expected utility of each element of the choice set in one single econometric model to predict behaviour. The dependent variable is a binary variable that corresponds to the alternative chosen by each subject. For example, if a subject chose to send \$20, the subject-alternative observation which corresponds to choosing \$20 takes the value of 1 and the rest of the subject-alternative observations (which the subject did not choose to send) take the value of 0. As main independent variables we include the following: the expected returns which are conditional on the quantity sent and vary by individual and within individuals, and the possible amount sent (\$5, \$10, \$15, ..., \$50). We run one single regression with 4,490 observations using a logit model with individual fixed effects.

Not surprisingly, sending higher amounts is less likely. More interestingly, panel (c) shows that a higher expected return significantly increases the likelihood of sending that quantity. Note that this result is stronger than those obtained in panels (a) and (b), since it controls for an individual fixed effect. Thus, not only individuals with higher expectations send more but also an individual is more likely to send more when she has particularly high expectations for a given amount sent.

3. Does the WVS-trust Question Capture Beliefs in Others' Trustworthiness?

Since the WVS-trust question aims to measure generalised trust, we expect this question to capture the expectation component of the trust game. However, this expectation does not need to be linearly proportional to the act of trust across the different possible amounts sent (in fact, it is not). As it appears above, for low amounts of money sent, beliefs do not enter the behaviour of senders; our interpretation is that, for these amounts, such behaviour may be closer to an act of charity than to an act of trust. However, as senders pass more money on to the receivers, they may hope to inspire acts of reciprocity. If this is true, it should be reflected in the correlation between WVS-trust question and beliefs: one would expect the WVS-trust measure to be correlated with beliefs only for quantities sent above \$12.50. Hence, not only is it important to determine whether the WVS-question is related to the expectation of trust at all but it is also important to establish to what level of expectations it is connected to. This is what we determine next.

Table 2
WVS-trust

	If sent \$5	If sent \$10	If sent \$15	If sent \$20	If sent \$25	If sent \$30	If sent \$35	If sent \$40	If sent \$45	If sent \$50
Expected return (conditional)	0.02 (0.03)	-0.0034 (0.01)	-0.0002 (0.01)	0.002 (0.008)	0.003 (0.007)	0.008 (0.006)	0.009* (0.005)	0.007* (0.004)	0.007* (0.004)	0.007*** (0.003)
Risk tolerance	0.01** (0.005)	0.01** (0.006)	0.01** (0.006)	0.01** (0.006)	0.01** (0.005)	0.01* (0.006)	0.01* (0.006)	0.01* (0.005)	0.01** (0.005)	0.01* (0.005)
Constant	-0.66 (0.47)	-0.54 (0.47)	-0.56 (0.48)	-0.60 (0.48)	-0.63 (0.48)	-0.7 (0.4)	-0.84* (0.48)	-0.82* (0.48)	-0.86* (0.49)	-0.96* (0.49)
Observations	434	434	434	434	434	434	434	434	434	434

Notes: Standard errors are in parentheses. **Significant at 1%, ***Significant at 5% and *at 10%.

In Table 2, we regress the WVS-trust question on the sender's expectation of money returned by the receiver.

For low amounts sent (less or equal to \$30), the expected return has a positive coefficient, but it is not economically or statistically significant. For higher amounts, however, the coefficient of expected returns is statistically and economically significant suggesting that the WVS-trust question captures beliefs in other people's trustworthiness: the odds that a subject responds that most people can be trusted increase by 0.7% with each dollar increase in expected returns.

Yet, the WVS-trust question does not seem to capture the pure expectation component since the variable risk tolerance has a positive and statistically significant effect. So risk aversion seems to play a role in determining the answer to the question whether 'most people can be trusted'.

Alternatively, we estimated (not reported) a different specification where we put the beliefs derived from the trust game on the left-hand side of the regression and the WVS-trust variable on the right-hand side.²¹ Fearing that the WVS-trust measure might be noisy, we instrument it with two other measures of trust we asked in our survey: the probability of recovering a lost wallet and students' trust in the University of Chicago. If the effect of the WVS-trust variable we found were spurious, it should disappear when we instrument. But if it is real, the estimated coefficient should increase because the IV estimate eliminates the attenuation bias typical of coefficients of noisy proxies. We find the latter to be true. According to our estimates, when subjects send \$50, those who respond that most people can be trusted expect back 50% more than subjects who disagree with this statement.²² Our results suggest that the WVS-question is a good, while noisy, proxy of the expectation component of the trust game.

4. Discussion

Having established that we are most interested in the belief component of the trust game, in this Section we investigate how these beliefs are formed and how they are correlated with the WVS-trust question. Our hypothesis is that in homogenous samples, like ours,²³ introspection might be the best way for subjects to predict their opponent's behaviour. If this is the case, a subject's beliefs should be correlated with her own trustworthiness. As Table 3 – panel (a) shows, this correlation is very strong in our sample.

The dependent variable is the amount a receiver returns to the sender, that is, her trustworthiness. This is highly correlated with what she expects to receive had she sent that quantity. The expectation alone explains between 33% and 38% of the cross-sectional variability in the amount returned. For every extra dollar, a subject

²¹ Results available from the authors.

²² Given that the instruments we use are variables collected in the same survey where the WVS-trust question was asked, concerns may arise regarding spurious correlation between these variables. Indeed, if a subject suffered from a given shock prior to responding these questions (e.g. her wallet was stolen), she may have responded in a similar way to all the questions causing spurious correlation among them. This can affect the first stage regression results, but should not affect the way subjects answered in the trust game since, as explained in the description of the games above, subjects completed the survey on average 7.33 days prior to playing the games.

²³ MBA students in our sample are of similar age range, life experience and background, all interested in studying business at Chicago, and with presumably similar interests and future objectives.

Table 3

Panel (a): Expected Return and Trustworthiness. Panel (b): Amount Sent and Trustworthiness. Panel (c): WVS-Trust and Trustworthiness

	If sent \$5	If sent \$10	If sent \$15	If sent \$20	If sent \$25	If sent \$30	If sent \$35	If sent \$40	If sent \$45	If sent \$50
<i>Panel (a)</i>										
Amount returned (conditional)	0.56*** (0.04)	0.60*** (0.03)	0.61*** (0.04)	0.54*** (0.04)	0.61*** (0.04)	0.58*** (0.04)	0.59*** (0.03)	0.61*** (0.03)	0.59*** (0.03)	0.59*** (0.04)
Constant	2.40*** (0.19)	5.25*** (0.36)	8.15*** (0.58)	11.98*** (0.80)	13.19*** (1.02)	17.94*** (1.27)	19.77*** (1.45)	21.07*** (1.67)	25.12*** (1.82)	28.67*** (2.19)
Observations	502	502	502	502	502	502	502	502	502	502
R ²	0.326	0.383	0.353	0.313	0.368	0.34	0.362	0.383	0.381	0.356
<i>Panel (b)</i>										
Amount returned (conditional)	-0.16 (0.19)	-0.04 (0.10)	0.1 (0.07)	0.15*** (0.06)	0.19*** (0.05)	0.20*** (0.04)	0.19*** (0.03)	0.17*** (0.03)	0.16*** (0.02)	0.15*** (0.02)
Risk tolerance	0.19*** (0.04)	0.19*** (0.04)	0.18*** (0.04)	0.18*** (0.04)	0.17*** (0.04)	0.17*** (0.04)	0.16*** (0.04)	0.17*** (0.04)	0.17*** (0.04)	0.16*** (0.04)
Other-regarding preferences	0.71** (0.33)	0.70** (0.33)	0.61* (0.33)	0.56* (0.33)	0.45 (0.33)	0.39 (0.33)	0.35 (0.32)	0.36 (0.32)	0.33 (0.32)	0.28 (0.32)
Constant	3.5 (3.28)	3.26 (3.28)	2.21 (3.28)	1.31 (3.27)	0.23 (3.26)	-1 (3.25)	-1.65 (3.22)	-1.89 (3.22)	-2.16 (3.21)	-2.15 (3.17)
Observations	478	478	478	478	478	478	478	478	478	478
R ²	0.055	0.054	0.058	0.068	0.084	0.102	0.118	0.12	0.128	0.141
<i>Panel (c)</i>										
Amount returned (conditional)	-0.014 (0.03)	-0.01 (0.01)	0.001 (0.01)	0.009 (0.01)	0.01 (0.01)	0.016*** (0.01)	0.014*** (0.01)	0.012*** (0.00)	0.013*** (0.00)	0.012*** (0.00)
Constant	0.430*** (0.14)	0.452*** (0.15)	0.365** (0.16)	0.214 (0.18)	0.134 (0.20)	-0.11 (0.20)	-0.145 (0.21)	-0.14 (0.20)	-0.243 (0.20)	-0.273 (0.20)
Observations	454	454	454	454	454	454	454	454	454	454

Notes. Standard errors are in parentheses. ***Significant at 1%, **at 5% and *at 10%.

expects a random opponent to return, she will increase the amount returned by 60 cents. This relationship is very stable for all the amounts involved.

Since beliefs are correlated with the amount returned, a subject's own trustworthiness should be correlated with her behaviour as a sender. Nevertheless, we have seen that at low quantities sent, other considerations are paramount. Therefore, we should find that trustworthiness is not predictive of the amount sent when it reflects the trustworthiness at low amounts sent but it should be predictive when trustworthiness is measured when high stakes are involved. We test this hypothesis in Table 3 – panel (b).

As expected, for amounts sent that are \$15 or lower, trustworthiness does not predict the amount sent while for higher amount sent, it does. This is true even after controlling for the other determinants of the amount sent, that is, risk aversion and other-regarding preferences.^{24, 25}

Finally, it is interesting to study the correlation between the WVS-trust question and trustworthiness. We know from Table 2 that the amount expected back is correlated with the WVS measure of trust. If beliefs are formed by taking into account a subject's own trustworthiness, then it is possible that trustworthiness is also correlated with WVS-trust. Table 3 – panel (c) shows that this pattern holds in our data.

It might seem disturbing that trustworthiness is correlated with the WVS-trust question since this question tries to measure trust in general. In fact, if people answer the WVS-question just extrapolating their own behaviour, their answer to that question might not be all that revealing. Yet, this correlation is found in other homogeneous samples like the Glaeser *et al.* (2000) sample but not in the Fehr *et al.* (2003) sample, in which subjects differ.^{26, 27}

One possible explanation for this puzzle is that the subjects' answer to the WVS-trust question may be influenced by the context in which this question is asked. For example, when the WVS-question is asked within an experiment conducted among Chicago MBAs, subjects might take other MBA students as the reference group, since these are the people they spend most of their time with. Since in such an experiment

²⁴ The trustworthiness of a subject (i.e. the amount returned in the trust game) may capture, in addition to the subject's own trustworthiness, his/her other-regarding preferences. This may be a reason why in Table 3 – panel (b), other-regarding preferences lose their significance when the trustworthiness of the subject starts being significant.

²⁵ The analysis of Table 3 – panels (a) and (b) can easily be reproduced in an econometric model that integrates all observations in the same regression with individual fixed effects (in the style of Table 1 – panel (c)). Such a model yields qualitatively the same results as Table 3 – panels (a) and (b), and is available from the authors.

²⁶ The Glaeser *et al.* (2000) subject pool consists of a homogeneous group of Harvard undergraduates (with self-selection to participate in laboratory experiments). In Fehr *et al.* (2003), the subject pool corresponds to a sample of German households, thus probably much more heterogeneous than the Harvard subject pool.

²⁷ An additional difference among the three papers we are examining concerns the initial endowment and multiplier of the amount sent. Glaeser *et al.* (2000) give \$15 as initial endowment only to the sender and the amount sent is multiplied by 2. Fehr *et al.* (2003) endow both senders and receivers with 10 euros and the amount sent is doubled as well as the amount returned. In our experimental design, we give \$50 to the sender as initial endowment and the amount sent is multiplied by 3. When only the sender is endowed, other-regarding preferences have an amplified effect. This is what we find in our article. But, other things being equal, this feature should not affect the relationship between the WVS-trust question and sender behaviour in the trust game. Moreover, when the quantity sent is multiplied by 2 instead of 3, there is a lower incentive to send if one expects nothing back. However, Fehr *et al.* (2003) still find evidence of trust in the sender behaviour.

subjects are similar to the reference group, it makes sense to base one's expectation of other people's trustworthiness on one's own trustworthiness. In contrast, when the question is asked in a much broader group, like the German population at large in Fehr *et al.* (2003), subjects might not extrapolate from their own behaviour because they might not be representative of the sample at large. Further tests are needed to validate this conjecture.

5. Conclusions

The concept of trust is playing an increasing role in economic analysis. This greater role requires an increased understanding of how to measure trust. Building on some widely used definitions, we argue that the sender's behaviour in a traditional trust game is a bad measure of trust. The trust game-based measure combines two components: the preferences of the sender and her beliefs in others' trustworthiness. Only the latter component is consistent with the prevailing definition of what trusts is. Ignoring this distinction leads to inconsistency in the measures of trust (Glaeser *et al.*, 2000).

If we accept that trust is the expectation about other people's behaviour, then both the answers to the WVS-question and the sender's expectation in a traditional trust game can be used as a measure.

Northwestern University, Kellogg School of Management, NBER and CEPR

Universidad Carlos III de Madrid

University of Chicago, Booth School of Business, NBER and CEPR

Submitted: 21 December 2009

Accepted: 12 December 2012

Additional Supporting Information may be found in the online version of this article:

Appendix A: Descriptive Statistics of the Sample.

Appendix B: Screenshots of the Trust Game.

Appendix C: Descriptive Statistics of Behavioural Data – Tables and Figures.

References

- Ashraf, N., Bohnet, I. and Piankov, N. (2006). 'Decomposing trust and trustworthiness', *Experimental Economics*, vol. 9(3), pp. 193–208.
- Bellemare, C. and Kroeger, S. (2007). 'On representative social capital', *European Economic Review*, vol. 51(1), pp. 183–202.
- Berg, J., Dickhaut, J. and McCabe, K. (1995). 'Trust, reciprocity and social history', *Games and Economic Behaviour*, vol. 10(1), pp. 122–42.
- Bloom, N., Sadun, R. and Van Reenen, J. (2009) 'The organisation of firms across countries', Working Paper. CEP Discussion Papers, 937. Centre for Economic Performance, London School of Economics and Political Science.
- Bohnet, I. and Zeckhauser, R. (2004). 'Trust, risk and betrayal', *Journal of Economic Behaviour & Organization*, vol. 55(4), pp. 467–84.
- Bohnet, I., Greig, F., Herrmann, B. and Zeckhauser, R. (2008). 'Betrayal aversion', *American Economic Review*, vol. 98(1), pp. 294–310.

- Bolton, G. and Ockenfels, A. (2000). 'ERC: A theory of equity, reciprocity, and competition', *American Economic Review*, vol. 90(1), pp. 166–93.
- Brandts, J. and Charness, G. (2000). 'Hot vs. cold: sequential responses and preference stability in experimental games', *Experimental Economics*, vol. 2(3), pp. 227–38.
- Brandts, J. and Charness, G. (2011). 'The strategy method versus the direct-response method: a first survey of experimental comparisons', *Experimental Economics*, vol. 3(3), pp. 375–98.
- Charness, G. and Rabin, M. (2002). 'Understanding social preferences with simple tests', *Quarterly Journal of Economics*, vol. 117(3), pp. 817–69.
- Cox, J.C. (2004). 'How to identify trust and reciprocity', *Games and Economic Behaviour*, vol. 46(2), pp. 260–81.
- Cox, J.C., Friedman, D. and Sadiraj, V. (2008). 'Revealed altruism', *Econometrica*, vol. 76(1), pp. 31–69.
- Falk, A. and Fischbacher, U. (2006). 'A theory of reciprocity', *Games and Economic Behaviour*, vol. 54(2), pp. 293–315.
- Fehr, E. and Schmidt, K.M. (1999). 'A theory of fairness, competition, and cooperation', *The Quarterly Journal of Economics*, vol. 114, pp. 817–68.
- Fehr, E. and Schmidt, K. (2006). 'The economics of fairness, reciprocity and altruism – experimental evidence and new theories', in (S.C. Kolm and J.M. Ythier, eds.), *Handbook of the Economics of Giving, Altruism and Reciprocity*, pp. 615–91, vol. 1, Amsterdam: Elsevier B.V.
- Fehr, E., Fischbacher, U., von Rosenblatt, B., Schupp, J. and Wagner, G. (2003). 'A nationwide laboratory. Examining trust and trustworthiness by integrating behavioural experiments into representative surveys', Working Paper, CEPR Discussion Papers, 3858.
- Fischbacher, U. (2007). 'Z-tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics*, vol. 10(2), pp. 171–8.
- Fischbacher, U., Gächter, S. and Fehr, E. (2001). 'Are people conditionally cooperative? Evidence from a public goods experiment', *Economic Letters*, vol. 71(3), pp. 397–404.
- Gambetta, D. (2000). 'Can we trust trust?', in (D. Gambetta, ed.), *Trust: Making and Breaking Cooperative Relations*, pp. 213–27, Oxford: University of Oxford.
- Glaeser, E., Laibson, D., Scheinkman, J.A. and Soutter, C.L. (2000). 'Measuring trust', *Quarterly Journal of Economics*, vol. 115(3), pp. 811–46.
- Guiso, L., Sapienza, P. and Zingales, L. (2004). 'The role of social capital in financial development', *American Economic Review*, vol. 94(3), pp. 526–56.
- Guiso, L., Sapienza, P. and Zingales, L. (2008). 'Trusting the stock market', *Journal of Finance*, vol. 63(6), pp. 2557–600.
- Guiso, L., Sapienza, P. and Zingales, L. (2009). 'Cultural biases in economic exchange', *Quarterly Journal of Economics*, vol. 124(3), pp. 1095–131.
- Holt, C.A. and Laury, S.K. (2002). 'Risk aversion and incentive effects', *American Economic Review*, vol. 92(5), pp. 1644–55.
- Houser, D., Schunk, D. and Winter, J. (2010). 'Distinguishing trust from risk: an anatomy of the investment game', *Journal of Economic Behaviour and Organization*, vol. 74(1–2), pp. 72–81.
- Isaac, R.M., Walker, J.M. and Thomas, S.H. (1984). 'Divergent evidence on free riding: an experimental examination of possible explanations', *Public Choice*, vol. 43(2), pp. 113–49.
- James, H.S., Jr. (2002). 'The trust paradox: a survey of economic inquiries into the nature of trust and trustworthiness', *Journal of Economic Behaviour and Organization*, vol. 47(3), p. 291–307.
- Karlan, D.S. (2005). 'Using experimental economics to measure social capital and predict financial decisions', *American Economic Review*, vol. 95(5), pp. 1688–99.
- Knack, S. and Keefer, P. (1997). 'Does social capital have an economic payoff? A cross-country investigation', *The Quarterly Journal of Economics*, vol. 112, pp. 1252–88.
- La Porta, R., Lopez de Silanes, F., Shleifer, A. and Vishny, R. (1997). 'Trust in large organisations', *American Economic Review*, vol. 87(2), pp. 333–8.
- Lazzarini, S.G., Madalozzo, R.C., Artes, R. and Siqueira, J.O. (2003). 'Measuring trust: an experiment in Brazil', Working Paper, WPE 01-2004, IBMEC Business School.
- Levine, D.K. (1998). 'Modelling altruism and spitefulness in experiments', *Review of Economic Dynamics*, vol. 1(3), pp. 593–622.
- Levitt, S.D. and List, J.A. (2007). 'What do laboratory experiments measuring social preferences tell us about the real world', *Journal of Economic Perspectives*, vol. 21(2), pp. 153–74.
- Manski, C.F. (2002). 'Identification of decision rules in experiments on simple games of proposal and response', *European Economic Review*, vol. 46(4–5), pp. 880–91.
- Marwell, G. and Ames, R.E. (1981). 'Economists free ride, does anyone else?', *Journal of Public Economics*, vol. 15(6), pp. 295–310.
- Offerman, T., Sonnemans, J., van de Kuilen, G. and Wakker, P.P. (2009). 'A truth-serum for non-bayesians: correcting proper scoring rules for risk attitudes', *Review of Economic Studies*, vol. 76(3), pp. 1461–89.
- Rabin, M. (1993). 'Incorporating fairness into game theory and economics', *American Economic Review*, vol. 83(5), pp. 1281–302.

- Reuben, E., Sapienza, P. and Zingales, L. (2008). 'A description of the Chicago-Templeton longitudinal study', Working Paper, University of Chicago.
- Schechter, L. (2007). 'Traditional trust measurement and the risk confound: an experiment in rural Paraguay', *Journal of Economic Behaviour and Organization*, vol. 62(2), pp. 272–92.
- Selten, R. (1967). 'Die strategiemethode zur erforschung des eingeschränkt rationalen verhaltens im rahmen eines oligopolexperiments', in (H. Sauermann ed.), *Beiträge zur Experimentellen Wirtschaftsforschung, Tübingen*, pp. 136–68. Tübingen: Mohr.
- Vyrastekova, J. and Onderstal, S. (2005). 'The trust game behind the veil of ignorance: a note on gender differences', CentER Discussion Paper 2005-96, Tilburg University.