TRUST IN TRIADS: AN EXPERIMENTAL STUDY†

Vincent Buskens[a,b], Werner Raub[a,c], and Joris van der Veer[a]

[a] Department of Sociology/ICS, Utrecht University, Netherlands
[b] Erasmus School of Law, Erasmus University Rotterdam, Netherlands
[c] Corresponding author: Department of Sociology/ICS, Utrecht University, Heidelberglaan 2, 3584 CS Utrecht, Netherlands, w.raub@uu.nl

Version May 2, 2010

Forthcoming in *Social Networks*

Abstract
Pairs of trustors play finitely repeated Trust Games with the same trustee in a laboratory experiment. We study trustfulness of the trustor and trustworthiness of the trustee. We distinguish between learning and control effects on behavior. Learning effects are related to the trustor's information on past behavior of the trustee. Control effects are related to the trustor's opportunities for sanctioning a trustee in future interactions. Hypotheses on learning and control effects are derived from backward-looking learning models and from forward-looking models of strategic behavior. The design of the experiment, with respect to trustfulness, allows for disentangling learning effects from a trustor's own experience with the trustee and learning effects through third-party information. Also, the design enables disentangling control effects on trustworthiness and trustfulness through a trustor's own sanction opportunities and opportunities for third-party sanctions. We find evidence for learning and control effects. The trustor's own experiences, the experiences of the other trustor, as well as the trustor's own sanction opportunities affect trustfulness. We find evidence for control effects on trustworthiness, including effects arising from opportunities for third-party sanctions. However, there is no evidence for control effects through opportunities for third-party sanctions on trustfulness. This could indicate limited strategic rationality of trustors.

Keywords: control, embeddedness, learning, trust

---

† We thank two anonymous reviewers, Davide Barrera, Chris Snijders, and Jeroen Weesie as well as seminar participants at Utrecht University and Venice International University for helpful comments and suggestions on earlier versions. This article is part of Buskens' project "Third-Party Effects in Cooperation Problems" funded by the Royal Netherlands Academy of Arts and Sciences (KNAW) and the Utrecht University High Potential-program of Vincent Buskens and Stephanie Rosenkranz "Dynamics of Cooperation, Networks, and Institutions." Raub acknowledges funding by the Netherlands Organization for Scientific Research (NWO) for the research projects "The Management of Matches" (S 96-168 and PGS 50-370) and "Commitments and Reciprocity" (400-05-089).

Trust in Triads: An Experimental Study

1 Introduction

Social and economic exchange often presupposes trust. When lending a book to a colleague, we are trustful that the colleague will return the book in good shape. A buyer of a second-hand car is trustful that the dealer is honest about hidden defects of the car. We may be more trustful when a colleague has often returned our books in good shape. Information from your friends on their good experiences with a second-hand car dealer may make you more inclined to buy a car from that dealer unless, maybe, you happen to know that the dealer is about to retire and close down his outlet. These examples illustrate that positive information about trustworthiness of the trustee in the past may foster trustfulness of the trustor. The final example illustrates that trustfulness and trustworthiness might become problematic when future sanctioning of the trustee's present behavior becomes infeasible.[1]

Our examples are related to the "embeddedness" of trust problems and exchange (Granovetter 1985). Embeddedness refers to repeated transactions over time between the same partners and to transactions between partners who share a network with third parties. Buskens and Raub (2002) specify two mechanisms through which embeddedness affects trust: learning and control. On the one hand, trustors can learn that a trustee has been trustworthy in the past and may infer from this that the trustee is likely to be trustworthy now as well. On the other hand, actors can base their trust on control through sanction opportunities in the future. The more opportunities a trustor has to sanction present behavior of the trustee in future interactions, the more likely the trustee is to realize that his short-term gains from abusing trust do not compensate for the long-term losses he would incur from sanctions of the trustor. Consequently, if the trustor has more control – that is, more opportunities for future sanctions of present trustee behavior – the trustee is more likely to be trustworthy and, hence, the trustor can be more trustful. Buskens and Raub (2002) also distinguish between two levels of embeddedness: the level of the dyad and that of the network. Learning and control operate at both levels (see Yamagishi and Yamagishi 1994: 138-139 for a similar discussion of learning and control through network embeddedness). Trustors can learn through their own experiences and through the experiences of others. Sanctions can be executed by the trustor herself or by third parties such as other trustors of the trustee. Note that sanctioning can be positive as well as negative. Future punishments for untrustworthy behavior today are negative sanctions. For example, we do not lend again a book to a colleague who today returned a book with coffee spots on several pages. But sanctions can also be positive and can involve rewards. For example, we inform our

[1] We use "trust" as shorthand for "trustfulness *and* trustworthiness." When referring specifically to trustfulness of the trustor, we write "the trustor places trust." When referring specifically to trustee behavior, we write "the trustee honors trust" (he is trustworthy) or "the trustee abuses trust."

friends about good experiences with a car dealer and thus produce new business opportunities for the dealer.

Previous experimental and survey research provides evidence for effects of embeddedness on trust (see Buskens and Raub 2010 for an overview). We improve in five ways on previous research. First, most of the available evidence suffers from the difficulty of disentangling the effects of learning and control. Only a few contributions (Rooks et al. 2000; Buskens and Weesie 2000; Batenburg et al. 2003; Rooks et al. 2006) focus on disentangling learning and control. These studies indicate that both learning and control affect trust. However, various network effects in these studies can still be interpreted as both learning effects and as control effects. Therefore, *how* embeddedness affects trust is a largely unresolved issue. Our experimental design allows us to more clearly distinguish between learning and control effects. In the experiment, pairs of trustors play finitely repeated Trust Games with the same trustee. The design of the experiment includes two conditions that represent different types of embeddedness of Trust Games. In one condition, each trustor only knows what happens in her own games with the trustee. In the other condition, each trustor also knows what happens in the games of the other trustor with the trustee. By tracing the decisions of subjects in the repeated Trust Game, our experiment allows us to disentangle learning and control effects at the network level. This will also allow for an assessment of the strength of different types of embeddedness effects.

Second, the interpretation of empirical findings on embeddedness effects in survey data is often problematic because embeddedness may be endogenous in that actors choose with whom to interact and with whom to exchange information (see Goyal 2007 and Jackson 2008 for overviews of the rapidly expanding literature on "strategic network formation" and Buskens 2002: 152-158 for a more specific discussion of problems in interpreting findings on embeddedness effects due to endogeneity of embeddedness characteristics). In contrast, the design of our experiment ensures that embeddedness characteristics are exogenous rather than being themselves results of individual choices.

Third, an additional problem of survey research on effects of network embeddedness is that third-party information can be problematic (Lorenz 1988; Raub and Weesie 1990: 648; Williamson 1996: 153-155). Why would a trustor be willing to provide information on trustee behavior to another trustor if providing such information is associated with costs in terms of, for example, time and effort? Moreover, trustors who are competitors may have incentives to provide misleading information about the trustee's behavior. Our design ensures that the provision of third-party information is not plagued by incentive problems and that such information is reliable.

Fourth, previous studies that do a rather good job at disentangling different kinds of embeddedness effects (Rooks et al. 2000; Buskens and Weesie 2000) employ vignette experiments and thus rely on subjects' responses to hypothetical situations without "real" incentives. The validity of such designs has been questioned (see Roth 1995: 5-8 for an overview of standard arguments on this issue). In our experiment, incentive compatibility is enhanced because the subjects' payments depend on their own and others' choices and they are precisely and truthfully informed about the incentive structure. This allows us to consider whether experimental findings are robust across design variations (see Buskens and Raub 2010 for a more general discussion of this issue) and, given robustness, we are able to attenuate external validity concerns.

Finally, previous empirical work studying embeddedness effects on trust typically focuses on testing hypotheses on trustfulness of the trustor, thereby neglecting tests of hypotheses on trustworthiness of the trustee (an exception that does focus on trustee behavior is Rooks et al. 2006). In the present study, we test hypotheses on the behavior of the trustor as well as the trustee.

In the following section, we conceptualize trust problems in a game-theoretic framework and specify hypotheses that follow from theoretical models on learning and control effects through dyadic and network embeddedness. Subsequently, we describe the details of the experiment and our analytical strategy. We analyze behavior of trustors and trustees in the results section. The final section offers conclusions and suggestions for further research.

2 Embedded Trust: Theories and Hypotheses

2.1 The One-Shot Trust Game

Building on Coleman's (1990: chapter 5) approach, we conceptualize a trust problem as a game with strategic interdependence between two actors, the trustor and the trustee. If the trustor places trust, the trustee has the possibility and an incentive to act opportunistically, that is, to abuse trust. Compared to the situation when she does not place trust, the trustor regrets when trust is abused but she is better off when the trustee honors trust. A trust problem can thus be modeled as the well-known Trust Game shown in Figure 1 (see also Camerer and Weigelt 1988; Dasgupta 1988; Kreps 1990a). The Trust Game starts with a move by the trustor, who chooses between placing trust (she is "trustful") and not placing trust. If the trustor does not place trust, the game is over, with trustor and trustee each obtaining a material payoff P (for example, points that are converted into money at the end of an experiment). If the trustor places trust, the trustee chooses between honoring trust (he is "trustworthy") and abusing trust. If the trustee honors trust, trustor and trustee each receive $R > P$. If the trustee abuses trust, the trustor receives $S < P$, while the trustee receives $T > R$.

FIGURE 1 ABOUT HERE

We start by analyzing the Trust Game under standard game-theoretic assumptions (see Rasmusen 2007 for a textbook): each actor is selfish in the sense of being exclusively interested in maximizing own payoffs ("utility = own money"), each actor acts rationally and expects rational behavior of the partner, and all features of the game, including the other actor's utility function, are known to both actors ("complete information"). Then, the trustee will abuse trust if the trustor places trust, because the payoff for abusing trust is larger than the payoff for honoring trust ($T > R$). The trustor realizes that the trustee would abuse trust. Thus, because the trustor is better off if she does not place trust than if she does place trust and trust is abused ($P > S$), she will not place trust. Bold lines in Figure 1 indicate these moves of trustee and trustor. In the technical language of game theory, not placing trust, while trust would be abused represents the unique subgame-perfect equilibrium of the Trust Game. Note that abuse of trust is not observed if trust is not placed. The Trust Game constitutes a social dilemma because both actors are worse off in the situation when trust is not placed than when trust is honored ($R > P$).

Remarks: (1) In addition to the unique subgame-perfect equilibrium the Trust Game has equilibria that are not subgame-perfect. These are strategy combinations such that the trustor does not place trust while the trustee would honor trust with probability $p \leq (P - S)/(R - S)$. Obviously, such equilibria are not subgame-perfect since the trustee, under our assumptions, would abuse trust with certainty after trust has been placed. (2) Note how the Trust Game relates to the Prisoner's Dilemma. While actors move simultaneously in the Prisoner's Dilemma, they move sequentially in the Trust Game. However, it would be misleading to consider the Trust Game simply as a sequential Prisoner's Dilemma. The Trust Game is best seen as a one-sided version of the Prisoner's Dilemma, since in the Trust Game there is only one actor who can profit from acting opportunistically while both actors have incentives for opportunism in the Prisoner's Dilemma. (3) It is clear from our discussion that we focus on trust as a result of incentive-driven behavior – "calculative trust" in Williamson's (1996) sense. We neglect trust in the sense of, for example, a general disposition to trustful behavior.

2.2 The Finitely Repeated Trust Game

To capture and distinguish embeddedness effects on trust theoretically as well as empirically we employ a model of finitely repeated Trust Games and an experimental design that closely matches this model (see Falk and Heckman 2009 for a recent account of the advantages of experimental designs in the social sciences). In the model and in our experiment, pairs of trustors are involved in Trust Games with the same trustee. More precisely, Trust Games are played in triads comprising two trustors and a trustee. Clearly, a triad represents a small network among the actors. First, one of the trustors, say, trustor 1, plays a Trust Game with the trustee. Next, the other trustor, trustor 2, plays a Trust Game with the same trustee. This pair of two games is played for a finite number of rounds. All actors, trustors and trustee, are informed about the whole structure of the repeated game such as the number of rounds to be played, each actor's payoff function, etc. Each actor also has full information on what happens in each game in which he or she is directly involved. After each Trust Game, depending on the experimental condition, information about what happened in that game might be transmitted to the other trustor who plays the next Trust Game with the trustee. All information, if sent, is transmitted immediately after a game has been completed. The theoretical assumption we use is that information, if sent to the other trustor, is truthful. In the experiment, information is sent automatically and truthfully by the system rather than actively by the trustors themselves. Therefore, the experimental design corresponds with the theoretical assumption. Both the theoretical model and the experiment employ two information conditions, with actors knowing in which information condition they are:

1. *No information exchange between trustors*: each trustor only knows what happens in her own Trust Games with the trustee but is not informed about what happens in the games of the other trustor playing with the same trustee.
2. *Full information exchange between trustors*: after each Trust Game, also the trustor not involved in that game receives information on the choices made in that game.

Under standard game-theoretic assumptions indicated above, a well-known result is that the finitely repeated Trust Game has a unique subgame-perfect equilibrium. The result follows from backward induction. In the final round, the trustee would abuse trust because this is his best option and there are no subsequent games that can be influenced by this choice. As a consequence, the trustor does not place trust in the final round. Therefore, the trustee would also abuse trust in the penultimate round. In this way, the argument iterates back to the first round and the trustor will never place trust in the subgame-perfect equilibrium. This result does not change if the trustee plays with two trustors as described above and it holds for both information conditions (again, while the subgame-perfect equilibrium is unique, the finitely repeated Trust Game has other equilibria that are not subgame-perfect, see Raub 1988).

Experiments show that there is considerable individual variation in how subjects make choices, including substantial percentages of trustful and trustworthy behavior even in a one-shot Trust Game (Snijders 1996; Snijders and Keren 2001; see Camerer 2003: chapter 2.7 for an overview). To keep things tractable theoretically and empirically, we focus on a widely used potential explanation for such variation, namely, that subjects differ in the way they value material payoffs due to preferences for fairness and inequity aversion (Fehr and Schmidt 1999; Bolton and Ockenfels 2000; see Camerer and Weigelt 1988; Dasgupta 1988 for formal models of finitely repeated Trust Games that capture such an approach). This means that actors' utilities may differ from their own material payoffs. Also, while knowing their own utility function, actors may be incompletely informed on the utility functions of other actors. In the case of the Trust Game, heterogeneity among actors with respect to preferences for fairness and inequity aversion combined with incomplete information on others' preferences boils down to the possibility that a trustee would honor trust in a one-shot Trust Game. This is so if he experiences "enough" guilt from inequity aversion after abusing trust. The trustor is uncertain about the trustee's utility function and only has beliefs about the probability that the trustee's inequity aversion is large enough. If that probability is large enough, the trustor has reason to place trust. Note that because the trustor moves first in the Trust Game, uncertainties of the trustee about incentives of the trustor do not directly affect trustee behavior in a Trust Game. Thus, in this paper, we do not address learning of the trustee but instead focus exclusively on the uncertainty of trustors about the trustee and how trustor behavior in the current round of the repeated Trust Game depends on experiences from earlier rounds of that repeated game.[2]

2.3 Backward-Looking Learning Models

Pure learning models are popular in sociological accounts of embeddedness effects on trust (Granovetter 1985: 400; Coleman 1990: chapter 5; Burt and Knez 1995). These are models of "backward-looking" adaptive behavior based on past experiences. Pure learning models assume that actors do not look ahead and, consequently, that they take into account neither the future sanction opportunities nor that others might anticipate on such sanction opportunities. Thus, these models focus exclusively on learning effects of embeddedness and neglect control effects.

[2] See Buskens and Raub (2010: 14-15) on how to make "enough guilt" and "high enough probability" precise and explicit. In some previous experiments, incomplete information is included in the experimental set-up by explicitly introducing two types of trustees, including one type that does not have a material incentive for abusing trust. Camerer and Weigelt (1988) provide some trustees with no incentive to abuse trust. Neral and Ochs (1992) let the type of trustee that has no incentive to abuse trust be played by the experimenter. Anderhub et al. (2002) let a robot play the trustee type without incentive to abuse trust, but have a control condition without this type of trustee. Brandts and Figueras (2003) force some trustees to honor trust. Earlier experiments show that some trustees who do not abuse trust can always be expected, even if this is not forced by experimental design. Then, artificially adding trustees who cannot abuse trust only changes the proportion of such trustees, but still does not provide perfect control over the expected proportion. Therefore, we will not include different types of trustees explicitly and by design in the experiment.

Different learning models for strategically interdependent actors have been developed and applied (see Fudenberg and Levine 1998 and Camerer 2003: chapter 6 for overviews). Two important variants are belief learning and reinforcement models. In belief learning, actors use their experience to update their beliefs about the expected behavior of their partners and choose a best response, given their beliefs. Two extreme forms of updating are Cournot (1838) dynamics and fictitious play as introduced by Brown (1951; see Camerer 2003: 283 on the name "fictitious play"). Following Cournot, actors play a best response against their opponents' last move: they believe that others will play the same as in the previous round, while older experiences are neglected. At the other extreme, in fictitious play, all past moves of the other actors are equally important and the relative frequencies of past moves are the beliefs about what other actors will do in the present round. It may seem unlikely that the trustee's moves prior to his last one do not matter at all, while it may also seem unrealistic that experiences from long ago are as important as recent experiences. Weighted fictitious play is a way of belief updating that includes decay in the importance of past experience (see Cheung and Friedman 1997; Camerer 2003: chapter 6). This approach is in line with Gautschi's (2000) finding that more recent experiences are more important than distant experiences in determining an actor's present behavior in social dilemmas.

Reinforcement models employ the payoffs an actor received in earlier rounds for predicting the actor's behavior in the current round. The higher the payoff obtained after a specific move, the more likely that the actor will use this move again. An early model is due to Bush and Mosteller (1955). More recently, economists and sociologists have refined this approach (see Roth and Erev 1995; Erev and Roth 1998; Macy and Flache 2002). Like belief learning models, reinforcement models come in different flavors. Reinforcement can be based on accumulated payoffs over time, but also on average payoffs over time for specific moves of the focal actor. Again, it seems plausible that distant experiences are less important in terms of reinforcement.[3]

The substantive predictions of different learning models hardly differ for the repeated Trust Game in our experiment. They typically imply that trustfulness is more likely if trust has been honored more frequently in earlier rounds of the repeated game and also, accounting for the reduced importance of the more distant past, more recently. On the other hand, trustfulness will become less likely after trust has been abused. Therefore, we use counts of events that happened in the past. In these counts, we include a parameter for how important information from earlier rounds of the repeated Trust Game is compared to information from more recent rounds. Similarly, we include third-party information: experiences of the other trustor are covariates in

[3] Combinations of belief and reinforcement learning have been developed as well (see Camerer and Ho 1999; Camerer et al. 2002).

the statistical analyses in addition to own experiences. The hypotheses that will be tested in this way are the following:

Hypothesis 1 (dyadic learning – trustor): The more a trustor's trustfulness has been honored in the past, the more likely it is that this trustor is trustful; the more a trustor's trustfulness has been abused in the past, the less likely it is that this trustor is trustful.

Hypothesis 2 (network learning – trustor): The more information a trustor has that trustfulness of the other trustor has been honored in the past, the more likely it is that she will be trustful herself; the more information a trustor has that trustfulness of the other trustor has been abused in the past, the less likely it is that she will be trustful herself.

Remark: In game theory, learning models are typically used to study whether equilibrium arises as the long-run outcome when less than fully rational actors interact over time in the same game and adjust their behavior based on experiences. Note that we do *not* ask whether actors playing a repeated Trust Game time and again learn to play an equilibrium of the repeated Trust Game due to such processes. Rather, we are interested in a more short-term process, namely, the trustor's learning about preferences of the trustee *within* a repeated Trust Game.

2.4 Forward-Looking Models

Backward-looking learning models do not take into account that actors strategically anticipate on the future behavior and the learning capacities of their partners. Models for games with incomplete information (Harsanyi 1967/68) assume rational forward-looking behavior with learning in the sense of Bayesian updating of beliefs. Consider a finitely repeated Trust Game with a trustee who, with some positive probability, does not have an incentive to abuse trust. Now, if the trustor places trust in some round of the repeated game that is not the final round, trust may be honored for one of two very different reasons. First, the trustee may have no incentive at all to abuse trust because his inequity aversion is large. Second, the trustee does have an incentive to abuse trust in the sense that his utility from abusing trust in the one-shot Trust Game is higher than his utility from honoring trust. Thus, his inequity aversion is small but the trustee follows an incentive for reputation building. The trustee knows that if he abuses trust, the trustor can infer for sure that the trustee's utility from abusing trust in the one-shot Trust Game is higher than his utility from honoring trust and will thus never place trust again in future rounds. On the other hand, if the trustee honors trust, the trustor remains uncertain about the trustee's incentives and may place trust again in the future. Conversely, the trustor can anticipate such behavior of the trustee and may therefore be inclined to place trust. In this game, the trustor can control the trustee in that placing trust in future rounds depends on honoring trust in the current round (the trustor's future sanction opportunities) and the trustor can learn about the incentives of the trustee from the trustee's behavior in previous rounds. The result is a subtle interplay of a trustor who tries to learn about and to control the trustee, taking the trustee's incentives for reputation building into account, and a trustee who balances the long-term effects of his reputation and the short-term incentives for abusing trust, taking into account that the trustor anticipates this balancing.

Camerer and Weigelt (1988; see also Dasgupta 1988; Bower et al. 1997) describe a sequential equilibrium (Kreps and Wilson 1982) for the standard finitely repeated Trust Game with a trustee and one trustor. The equilibrium comprises placing and honoring trust in some – and possibly many – rounds. In this equilibrium, the trustor places trust and the trustee honors trust in early rounds of the repeated game. The trustor and a trustee with incentives for abusing trust start to randomize towards the end of the repeated game. For the trustee, the reason is that the negative sanction related to abuse of trust decreases and at some point is no longer large enough anymore to prevent him from reaping the short-term gains of abusing trust. The reason for the trustor is that she also realizes that the trustee may start to abuse trust and, therefore, she might stop placing trust before the trustee starts to abuse trust. As soon as there is one abuse of trust, or the trustor does not place trust, the trustor will never place trust in future rounds. Buskens (2003) analyzes such a model for the triads considered in our model and in our experiment. The equilibrium for the triads is similar to the equilibrium in the two-person game. The equilibrium remains the same when trustors do not exchange information about the behavior of the trustee. In the condition with full information exchange between trustors, actors start to randomize in equilibrium in a later round than in the condition with no information exchange. Trustor 2 can profit less from this network information at the end of the game because there is one game less to be played after her game as compared to trustor 1's game in the same round. This leads to the following hypotheses on control effects on trustworthiness and trustfulness.

Hypothesis 3 (dyadic control – trustee): The more rounds left in the game, the higher the likelihood that trust will be honored; the likelihood of trustworthiness decreases faster in the last few rounds of the game than in earlier rounds (end-game effect).

Hypothesis 4 (network control – trustee): In the condition with full information exchange between trustors, compared to the condition with no information exchange, the likelihood of trustworthiness is higher and will decrease less in early rounds. The end-game effect will be stronger in games with trustor 2 than in games with trustor 1.

Hypothesis 5 (dyadic control – trustor): The more rounds left in the game, the higher the likelihood that a trustor is trustful; the likelihood of trustfulness decreases faster in the last few rounds of the game than in earlier rounds (end-game effect).

Hypothesis 6 (network control – trustor): In the condition with full information exchange between trustors, compared to the condition with no information exchange, the likelihood of trustfulness is higher and will decrease less in early rounds. The end-game effect will be stronger for trustor 2 than for trustor 1.

Remarks: (1) Note that the hypotheses on control effects of embeddedness on trust are hypotheses on effects of future sanction *opportunities* of trustors rather than hypotheses on the *implementation* of (negative) sanctions by trustors. In fact, the game-theoretic model implies that increasing future sanction *opportunities* will increase trustworthiness of the trustee and therefore also trustfulness of the trustors so that there is less need for *implementing* negative sanctions. (2) One might ask whether some of the hypotheses on learning effects that have been derived from backward-looking models also follow from or are at least consistent with the forward-looking game-theoretic model and fully rational equilibrium behavior. Roughly, the forward-looking model likewise predicts that abuse of trust reduces the likelihood of future trustfulness. However, the forward-looking model does not predict that honored trust in early rounds of the repeated game increases the likelihood of trustfulness. The reason for this is that, in the sequential equilibrium, trustworthy behavior of the trustee allows the trustor to update her belief only if it occurs in later rounds since both trustees with and trustees without incentives to abuse trust in the one-shot Trust Game will honor trust for sure in early rounds of the repeated Trust Game (see Buskens 2003: 246-247 for a more detailed discussion and for additional predictions from the game-theoretic model that are not tested in this paper).

3 The Experiment

3.1 Procedure

In the experiment, the outcomes of the Trust Games are points that subjects earn. Not placing trust yields 10 points for the trustor as well as for the trustee; when trust is placed and honored, each actor receives 20 points; and when trust is abused, the trustee receives 40 points, leaving the trustor with no points (see the right-hand Trust Game in Figure 1). Subjects are paid 1 eurocent for each point they earn at the end of the experiment.

The experiment was computerized and programmed using z-Tree software (Fischbacher 2007). Subjects play repeated Trust Games with 15 rounds. Subjects are matched in groups of three for these 15 rounds, one trustee and two trustors, which we call triads. In each of the 15 rounds, the trustee plays one Trust Game with each of the two trustors. Therefore, in every round, while each trustor plays one Trust Game, the trustee plays two Trust Games, adding to 30 Trust Games played per repeated game by the trustee. The trustee does not necessarily need to make a choice in all these 30 games: when the trustor does not place trust, the trustee has no choice to make.

In every round, the trustee always plays with the same trustor first, while the other trustor has to wait, and always plays second with this trustee. Thus, within a repeated game, the trustors always move in the same order. The trustors are referred to as trustor 1 and trustor 2, respectively.

There are two conditions in the experiment. In both conditions, the trustee is immediately informed on the trustor's move in the current Trust Game. Between conditions, the amount of information is varied that is shared among the two trustors playing with the same trustee. In the "no information exchange between trustors" condition, trustors do not share any information. In this condition, subjects only know what happens in their own games with the trustee. So each trustor is immediately informed about the choices of the trustee in her own games with the trustee, but not about the choices of the other trustor and the trustee in their respective games. In the "full information exchange between trustors" condition, trustors playing with the same trustee do share all information about each other's games. Besides receiving information about the choices the trustee made in her own games with the trustee, a trustor receives information about the choices of the other trustor and the response of the trustee to those choices. In this condition, as soon as a game has been played, either the first or the second game in the round, also the other trustor receives information on the choices made in this game. Information is provided automatically by the computer and is always truthful. All subjects know what information condition they are in, and thus they also know what information is available to the other two actors in their triad. All subjects see the outcomes of the games they played themselves in previous rounds on their own screen. In the condition with full information exchange, each trustor also sees the outcomes of the previous Trust Games of the other trustor on the screen. All information is provided in a so-called history box at the bottom of the screen. This allows the actors to see and review how their partner(s) acted in past rounds. Note that the experimental design ensures that embeddedness is exogenous and is not itself a result of subjects' choices.

Every subject played three repeated Trust Games as described above, once as a trustee, once as trustor 1, and once as trustor 2. Each subject played all three repeated games in the same information condition. In between the three repeated games, the subjects were rematched to other subjects. Subjects were never rematched to other subjects they already played with in a previous repeated game. This procedure was made known to all subjects.

The experiment was conducted in the ELSE laboratory of Utrecht University, in a computer room specially designed for experiments: every subject was seated behind a PC in his or her own cubicle, with a separate cubicle for the experimenters. Each session comprised 18 subjects. When all 18 subjects were present, instructions were provided on paper. Within a session, the instructions were the same for everyone and this was also made known to the subjects. Subjects were informed of how the game they were about to play worked, and that they would receive 1 eurocent for every point they earned.

After reading the instructions, a few questions were asked on the computer screen so that subjects could check whether they understood the instructions. In the case of incorrect answers, subjects were provided the correct answers and a brief explanation. Subjects also played two practice rounds in which they could not earn any points. These rounds were not played with other subjects, but with the computer to guarantee standardized moves of the "partner" in these rounds.

Then the three repeated Trust Games were played. At the end of the experiment, subjects had to fill in a short questionnaire, including items concerning their general trustfulness and trustworthiness. Meanwhile, the experimenters prepared closed envelopes with the earnings for each subject. The sessions took between 55 and 70 minutes. Subjects earned on average €10.67, with an average of €11.10 for subjects in the condition with full information exchange between trustors and €10.25 for those in the condition without information exchange. The minimum and maximum earnings were €7 and €12.40, respectively.

The information about the structure of the experiment such as the number of rounds, roles (trustor or trustee), and what subjects would get to know was honestly provided and subjects were never deceived or placed in any situation other than the one they were told. In order to prevent inducing normative associations, the names of the different actor roles and their possible moves were rendered neutrally. For instance, the moves of the trustee were labeled "down" and "right" rather than "honor trust" and "abuse trust."

3.2 Subjects

In total, 72 subjects participated in the experiment, with 28 male and 44 female participants. Subjects were mostly undergraduate students from different fields, most of them students of social sciences. Subjects were recruited using the online recruitment system ORSEE (Greiner 2004). Four sessions were scheduled and 18 subjects participated in each session. Two sessions were played in the condition with no information exchange between trustors and two sessions in the condition with full information exchange.

4 Data and Statistical Model

4.1 Data

The experiment comprised four sessions, with six triads per session and three repeated games of fifteen rounds per subject, each round comprising two Trust Games. Thus, $4\times6\times3\times15\times2 = 2160$ Trust Games were played in total. Clearly, trustee behavior is observed only in games in which the trustor places trust. There are 485 games in which the trustor did not place trust, leaving 1675 games (78% of the total number of games played) in which the trustee's behavior is observed. The trustor places trust more often in the condition with full information exchange between trustors (913 of the 1080 games) than in the other condition (762 of the 1080 games) and trustfulness decreases over rounds. Also, trustworthiness decreases over the rounds. Table 1 summarizes the percentages of trustfulness and trustworthiness per round. Figure 2 displays the same descriptives graphically.

TABLE 1 ABOUT HERE

FIGURE 2 ABOUT HERE

The descriptive findings are in line with our hypotheses and with earlier experiments (see Camerer 2003: 446-453; Buskens and Raub 2010 for overviews). Trustfulness and trustworthiness are high for most of the rounds, with strong end-game effects in the last couple of rounds. We observe that repeating the Trust Game with the same partner (dyadic embeddedness, also known as "partner matching" in the experimental literature), and the availability of third-party information (network embeddedness) have complementary effects. This could not be concluded from otherwise closely related experiments (Bolton et al. 2004; Bolton and Ockenfels 2009). In these earlier experiments there were conditions with either dyadic or network embeddedness but there was no condition with both dyadic *and* network embeddedness. Also, these earlier studies do not distinguish between learning and control in their analysis of embeddedness effects. To the best of our knowledge, Barrera and Buskens (2009) present the only laboratory experiment in which both dyadic and network embeddedness were manipulated in Trust Games in ways similar the present study. Their experiment, however, involves a more complex setting in which trustors play simultaneously rather than sequentially with one trustee. This affects the predictions, because sanction opportunities do not differ between the trustors playing with the same trustee.

In summary, the descriptive analyses show that dyadic embeddedness leads to rather high levels of trustfulness and trustworthiness, levels higher than those normally found in one-shot Trust Games, including series of one-shot Trust Games, each played with a different partner (also known as "stranger matching" in the experimental literature). Still, information exchange between trustors who play Trust Games with the same trustee about behavior of this trustee further enhances trustfulness and trustworthiness.

4.2 Variables

To test the hypotheses, we cannot restrict ourselves to comparing the two experimental conditions. Rather, we need to estimate the likelihood of placing and honoring trust depending on the experimental condition, past experiences of subjects, the number of rounds to be played, and whether trustor 1 or trustor 2 is playing. Our dependent variables for these analyses identify the behavior of the trustor and the trustee. The dichotomous variable TRUSTFULNESS equals 1 if the trustor places trust and 0 otherwise, while the dichotomous variable TRUSTWORTHINESS equals 1 if the trustee indeed honored trust after trust has been placed and 0 otherwise.

The dichotomous variable FULL INFORMATION equals 1 for the experimental condition with full information exchange between trustors and 0 for the other experimental condition.

Various independent variables represent past experiences. First, for each round t, we construct the weighted number of times a trustee honored trust in the past as

$$\text{PASTHONOR\_OWN}_t = \sum_{\tau=1}^{t-1} w_1^{t-\tau-1} \text{HONOR\_OWN}_\tau . \qquad (1)$$

where HONOR_OWN indicates whether trust placed by the focal trustor is honored in a specific round. We define PASTHONOR_OWN = 0 for the first round of a repeated game. Similarly, we define a variable for the number of times trust was abused in the past:

$$\text{PASTABUSE\_OWN}_t = \sum_{\tau=1}^{t-1} w_1^{t-\tau-1} \text{ABUSE\_OWN}_\tau . \qquad (2)$$

where ABUSE_OWN indicates whether trust placed by the focal trustor is abused in a specific round. Note that ABUSE_OWN and HONOR_OWN can be equal to 0 simultaneously, namely, when the trustor does not place trust. Therefore, the effects of the two variables defined above should be interpreted relative to the number of times the trustor placed trust. The parameter w_1

is chosen to be 0.5 in the statistical model (see below) implying that the experience in the previous round has a weight 1, the round before that a weight 0.5, the round before that a weight 0.25, and so on. We assume that w_1 is the same for honored trust and for abused trust. We will compare model fit for analyses with other values for w_1.

In the condition with full information exchange between trustors, each trustor also receives information about the games of the other trustor. To complete the set of independent variables needed for testing hypotheses on learning effects, we thus define variables for such third-party information:

$$\begin{aligned} \text{PASTHONOR\_OTHER}_t &= \sum_{\tau=1}^{t-1} w_2^{t-\tau-1} \text{HONOR\_OTHER}_\tau; \\ \text{PASTABUSE\_OTHER}_t &= \sum_{\tau=1}^{t-1} w_2^{t-\tau-1} \text{ABUSE\_OTHER}_\tau. \end{aligned} \tag{3}$$

where HONOR_OTHER and ABUSE_OTHER indicate whether trust placed by the other trustor in the same triad as the focal trustor is honored or abused in a specific round. The variables PASTHONOR_OTHER and PASTABUSE_OTHER are always set to 0 in the condition without information exchange between trustors. Information from the other trustor might be forgotten or discarded faster than own experiences. Therefore, we introduce a new parameter w_2. In the analyses, however, this parameter will also be set to 0.5, with additional robustness analyses similar to those for w_1.

The variables representing control effects are quite straightforward. FUTURE is the number of rounds left. FUTUREFULL is the interaction of FUTURE with a dummy for whether subjects are in the condition with full information exchange between trustors. This variable represents the additional effect of the number of rounds left in the condition with full information exchange. We use dummies that indicate the penultimate and the final round of the repeated game: ROUND14, ROUND15. These variables are again interacted with dummies for the information condition: ROUND14FULL, ROUND15FULL. We also distinguish between the games with trustor 1 and with trustor 2 in the last two rounds for the condition with full information exchange between trustors. Therefore, two dummies are used for the games with trustor 2 in rounds 14 and 15 (ROUND14TR2, ROUND14TR2). Another two dummies with three-way interactions indicate whether trustor 2 plays these last two rounds in the condition with full information exchange between trustors or in the other experimental condition (ROUND14TR2FULL, ROUND15TR2FULL). The variables for the control effects are the same in the analyses of TRUSTFULNESS and TRUSTWORTHINESS.

4.3 Statistical Model

The statistical model used to analyze TRUSTFULNESS is a three-level logistic regression model. Based on the difference in attractiveness between the two possible moves of the trustor and assuming that there is some randomness in the extent that we know the attractiveness, we can estimate a logistic regression model for the probability that trustor *j* in triad *i* places trust in round *t*:

$$P_{ijt}^{Trustful} = \frac{e^{\beta_1 \cdot (\text{dyadic learning}) + \beta_2 \cdot (\text{network learning}) + \beta_3 (\text{dyadic control}) + \beta_4 (\text{network control}) + u_i + v_{ij} + \varepsilon_{ijt}}}{1 + e^{\beta_1 \cdot (\text{dyadic learning}) + \beta_2 \cdot (\text{network learning}) + \beta_3 (\text{dyadic control}) + \beta_4 (\text{network control}) + u_i + v_{ij} + \varepsilon_{ijt}}}, \quad (4)$$

where the β's indicate the vectors of regression coefficients for the respective groups of independent variables, namely, variables representing learning from own experiences and learning through third-party information as well as variables representing control through own future sanction opportunities and through future sanction opportunities involving third parties. Furthermore, u_i is a random component for the triad in which the decision is made, v_{ij} is a random component for the trustor within the triad who makes the decision, and ε_{ijt} is the random component for each individual decision. This is a hierarchical three-level model (see, for example, Snijders and Bosker 1999). Strictly speaking, we have a cross-classified nesting because trustors are involved twice in a series of 15 Trust Games with different partners. However, estimating this more complex structure only marginally affects the outcomes of the analyses. Also, the random component related to trustors has a similar size in this more complex estimation. Another way of controlling for this interdependence is to use a fixed effect for whether a decision is taken as trustor 1 or trustor 2. That fixed effect is estimated to be 0. Because none of our results are affected by whether or not we take this additional nesting into account, we stick to the simpler hierarchical model. One remaining problem with the estimation is that, in principle, we would prefer a simultaneous estimation of how fast past experiences are forgotten (w_1 and w_2). Unfortunately, we were not able to simultaneously estimate these parameters and include the complex multilevel structure specified above. Therefore, we first estimated w_1 and w_2 simultaneously with the regression coefficients in a model that does not account for the multilevel structure but only employs robust standard errors corrected for clustering in triads (see Huber 1967; Rogers 1993). This analysis yields w_1 and w_2 both having a value close to 0.5 (where the value for w_1 is significant at the 0.001 level and the value for w_2 is marginally significant, indicating that w_2 could not be estimated very precisely). Subsequently, assuming that the values are actually both 0.5, we run the logistic regression with the same robust standard errors. These standard errors were virtually the same as in the simultaneous estimation. This ensures that we can also use the multilevel model specified above assuming that the parameters equal 0.5, without concern that the separate estimation of these parameters severely compromises our standard errors.

The statistical model for TRUSTWORTHINESS of trustee i towards trustor j in round t is similar:

$$P_{ijt}^{Trustworthy} = \frac{e^{\beta_1 (\text{dyadic control}) + \beta_2 (\text{network control}) + u_i + v_{ij} + \varepsilon_{ijt}}}{1 + e^{\beta_1 (\text{dyadic control}) + \beta_2 (\text{network control}) + u_i + v_{ij} + \varepsilon_{ijt}}}, \quad (5)$$

where the β's again represent the regression coefficients; u_i is a random component for the subject (trustee) making the decision. In this case, this is equivalent to the triad in which the decision is made; v_{ij} is a random component for the trustor in the triad with whom the trustee is playing a specific Trust Game; and ε_{ijt} is the random component for each individual decision. Again, the specification of the random components could have been more complex, because each trustor is involved in two triads. Therefore, the random component for trustors could have been specified as a cross-classified model in which the random component represents randomness related to a specific subject playing as a trustor. Because random components related to the trustors are consistently estimated to be 0 in the models for explaining trustees' behavior, we stick to the simpler hierarchical model in which we control for nesting of trustees' decisions within trustors.

We choose two different ways to cope with the first round in each repeated game. In the first round, trustors do not have any learning experiences yet. Moreover, we do not know the extent to which trustors expect trustees to care about sanction opportunities of the trustors. Due to control effects, we expect trustfulness to decline over time, but we do not know how high trustfulness will be in the first round. In addition, trustfulness in later repeated games might be affected by experiences in earlier repeated games. Therefore, we include a dummy for the first round for each repeated game in the models to predict TRUSTFULNESS. Ideally, we would have done the same for the trustees, but in the third repeated game all trustees who were trusted also honored trust. This led to estimation problems and, therefore, we excluded the first round observations from the model for trustees. Except for obtaining a very high coefficient for the first round in the third repeated game, including the first round cases and related dummies does not affect the substantive results of the analysis of trustee behavior. After excluding first round observations, we have 1542 cases left in the analyses of TRUSTWORTHINESS below.

It is clear from the set-up above that our main analyses are not straightforward experimental contrasts, but more complex statistical analyses. Figure 2 already shows that the main contrast predicted for the experimental conditions is present in the data. However, the hypothesized mechanisms that might cause the difference between the conditions cannot be distinguished merely by considering this contrast. Rather, we have to learn more about how subjects adapt and change their behavior during the experiment. Because the effects we discuss below are not exclusively based on experimental manipulations, we are, in principle, vulnerable to objections against our findings based on alternative mechanisms for which we do not control in the analyses. Still, we can exclude at least some alternative explanations. For example, because we have manipulated the network exogenously, we can exclude that patterns in the data are due to endogenous changes in the network (cf. Buskens 2002: chapter 5). Also, due to the longitudinal character of the data, we are sure that the learning effects we analyze are based on experiences *before* the focal interaction, which provides a major advantage over cross-sectional studies.

5 Results

5.1 Trustfulness

We first analyze TRUSTFULNESS. As a baseline model, we report a logistic regression with a dummy for the experimental condition (FULL INFORMATION), a dummy for whether trustor 2 is involved rather than trustor 1 (TRUSTOR2), and two dummies for the second (REPEATEDGAME2) and the third repeated game (REPEATEDGAME3) in Table 2. This model clearly shows that TRUSTFULNESS is higher in the condition with full information exchange between trustors than it is in the condition with no information exchange. TRUSTFULNESS also increases over the repeated games, as indicated by the significant difference between the first and the third repeated game, while the second repeated game is in between. The variance of the random part at the decision level ε_{ijt} equals 3.29 in this model. We thus conclude from the random parts at the triad and trustor level that about 28% of the variance can be attributed to the triad, while only 5% can be attributed to a specific trustor within a triad. By analyzing each condition separately, we found that almost all the variance at the trustor level is due to the condition with no information exchange, while there is hardly any variance that can be attributed to the individual trustors in the condition with full information exchange. This suggests that it does not matter much that we neglect individual differences between trustors, such as in individual tendencies to place trust: these individual differences can hardly be expected to explain a lot of variance and the differences seemed to be washed away anyway in the condition with full information exchange between trustors.

TABLE 2 ABOUT HERE

In the full model, the main effect of FULL INFORMATION vanishes, indicating that the difference between the experimental conditions is mainly due to the variables representing learning and control effects in the full model. The full model provides clear evidence for Hypothesis 1 on effects of dyadic learning and Hypothesis 2 on network learning. Trustors are more likely to place trust after experiencing more honored trust themselves and they are less likely to place trust after experiencing more abused trust themselves. In addition, when they observe more often that trust placed by the other trustor is honored, this also increases TRUSTFULNESS, while TRUSTFULNESS decreases if trustors observe more abused trust of the other trustor in the same triad. It can be seen that the effect of experiencing honored trust in the trustor's own games is larger than the effect of information on honored trust in games of the trustee with the other trustor. It is striking that the effect of information on abused trust in games of the trustee with the other trustor is almost as large as the effect of experiencing abused trust in the trustor's own games with the trustee. The significance of the effect of information on abused trust in games of the trustee with the other trustor is smaller as can be inferred from the larger standard error, but that could be due to the limited amount of data on these experiences because they only occur in the condition with full information exchange between trustors.

Considering control effects on TRUSTFULNESS, we see a clear dyadic control effect of the rounds still to be played. Also, there are end-game effects especially in round 15, while dummies for earlier rounds did not add to the explained variance. These results provide support for Hypothesis 5. However, there is not much support for Hypothesis 6 on effects of network control on TRUSTFULNESS. This is evident, first, from the fact that there is no effect of FULL INFORMATION in the full model. Also, the interaction of the number of rounds left with the full information exchange condition is not significant, indicating that the general decrease of TRUSTFULNESS is not attenuated in the condition with full information exchange between trustors. The end-game effects are about twice as strong in the condition with full information exchange than they are in the condition with no information exchange and this difference is significant for round 15. However, the end-game effects do not start later in the condition with full information exchange. Therefore, it might very well be that the steeper network effects are mainly due to earlier experiences of honored trust in the condition with full information exchange through which the level of trustfulness is higher before it starts to decrease. Clearly, this interpretation differs from the control interpretation. If the control interpretation were correct, one would expect a difference in trustfulness between trustor 1 and trustor 2, given that trustor 1 has more future sanction opportunities left in round 14 and round 15 than trustor 2. The additional dummies that interact the variables indicating the two final rounds with trustor 2 and the three-way interactions in which the condition with full information exchange is distinguished are not significant. The directions of the effects suggest that the end-game effects are slightly stronger for trustor 2, but mainly in the condition without information exchange between trustors. Since all four effects are far from significant, we do not find evidence for the corresponding part of Hypothesis 6.

By comparing separate models in which we only include either learning or control effects (for example, by comparing the loglikelihoods of the respective models), we could confirm that learning effects explain more of the trustfulness of trustors than control effects. Finally, the controls for the first rounds show that the starting level of trustfulness slowly increases over the repeated games, indicating that trustors increasingly realize that trustfulness can be beneficial in the beginning of a repeated game.

We implemented additional analyses to check for the robustness of the findings. First, we analyzed each of the repeated games separately. Learning effects are found even in these smaller datasets. The significance of the control effects is less stable due to the reduction of power, but the direction of the effects remains consistent over all analyses. Second, we implemented analyses with variables representing past experiences constructed with the two most extreme values for w_1 and w_2, namely both being equal either to 0 or 1. Here, 0 means that trustors only take the experiences from the previous round into account and 1 means that all past experiences are equally important. The analysis in which both parameters equal 0 provides very similar results, but the loglikelihood of this model (-634.610) is slightly worse than for the model presented in Table 2, indicating that the effects of past experiences are less well covered if w_1 and w_2 equal 0. The analysis with both parameters equal to 1 yields different results. First, this model fits much worse than our main model (loglikelihood = -698.948). Also, in this model, FUTURE has a negative effect. This can be explained if indeed experiences that happened far in the past were afforded too much weight. Because trustors are predominantly trustful in early rounds, this leads to an overestimation of the value of positive experiences later in a repeated game. The negative effect of FUTURE reflects this overestimation.

5.2 Trustworthiness

We now turn to the analysis of TRUSTWORTHINESS. The baseline model in Table 3 shows that there is more TRUSTWORTHINESS in the condition with full information exchange between trustors than in the condition without information exchange. While trustors develop some more trustfulness over the repeated games, the trustees' behavior does not change significantly over the repeated games. Controlling for these repeated game variations, we see that no variance is attributed to the trustor level and that 28% of the variance is attributed to the trustee level. The remaining 72% of the variance is at the decision level.

TABLE 3 ABOUT HERE

The full model in Table 3 shows the results of the tests of hypotheses on TRUSTWORTHINESS. First, notice that the difference between the two experimental conditions is not explained away by the hypothesized effects for the trustee. The coefficients in the models are scaled with the variance and therefore we cannot directly compare the two coefficients. However, the significance level of the dummy for the information condition remains similar. We provide a more substantive interpretation for this result below.

With respect to control effects on TRUSTWORTHINESS, the effect of dyadic embeddedness (FUTURE) is strongly significant. However, the interaction term with the condition with full information exchange shows that dyadic control is only present in the condition without information exchange between trustors. Apparently, control is so strong in the condition with information exchange between trustors (which is also indicated by the remaining main effect of FULL INFORMATION) that the likelihood of honoring trust remains at or even above the 90%-level throughout the first thirteen rounds of a repeated game. The fact that the main effect of FULL INFORMATION is positive and that TRUSTWORTHINESS does not decrease in the first 13 rounds in the condition with full information exchange thus indicates that there is an additional control effect of network embeddedness over and above dyadic control. After round 13, there is a clear drop in TRUSTWORTHINESS in both information conditions. When we study these end-game effects in more detail, we see that the trustee is less trustworthy with trustor 2 than with trustor 1 in the condition with full information exchange. Most notably that can be seen from the significant three-way interaction of the information condition with trustor 2 in round 15. This indicates an additional network control effect, because trustor 2, compared to trustor 1, has less (or no) sanction opportunities especially in these last two rounds. Summarizing, the results provide substantial support for Hypothesis 3 as well as for Hypothesis 4.

To conclude, we analyzed TRUSTWORTHINESS also for each of the repeated games separately. End-game effects become unstable in these additional analyses, because there are only few data points to estimate these effects. In two of the three analyses, the effects of the experimental condition as well as the effect of the number of rounds still to be played are reproduced. Also the interaction effect between the experimental condition and the number of rounds to be played is similar. This again indicates the robustness of the evidence for Hypotheses 3 and 4.

6 Conclusion and Discussion

In this paper, we have discussed an experiment in which pairs of trustors play Trust Games with the same trustee. This is the simplest set-up for simultaneously studying effects of dyadic embeddedness and network embeddedness on trust. We distinguished between learning and control effects of embeddedness. We have analyzed how both trustfulness and trustworthiness are affected by embeddedness.

For trustees, we find both dyadic control effects as well as network control effects on trustworthiness. Learning effects at the dyadic and the network level are both strong determinants of trustfulness of trustors. Dyadic learning effects on trustfulness seem slightly stronger than network learning effects. While we find dyadic control effects on trustfulness, we do not find evidence for network control effects on trustfulness. The higher levels of trustfulness under network embeddedness are actually due to the trustees who seem to anticipate the stronger sanction opportunities of the trustors. The trustees are more trustworthy under network embeddedness, thus implying that the trustors have more positive learning experiences that, in turn, lead to more trustfulness. Thus, network embeddedness can and, in our experiment, empirically does affect trustfulness even without direct network control effects on trustor behavior. Since trustees react to trustors' sanction opportunities through network embeddedness, network embeddedness increases trustworthiness. Through learning effects on trustor behavior, network embeddedness then also increases trustfulness.

Compared to earlier research, our experimental design has several strengths. The design allows us to clearly distinguish between learning and control effects of embeddedness as well as to better assess the relative strength of different embeddedness effects. Also, our design ensures that embeddedness characteristics are exogenously given rather than being themselves results of individual choices. Moreover, the design implements the theoretical assumption that third-party information, if provided at all, is reliably provided. These are advantages compared with survey research on how embeddedness affects trust. Furthermore, our design ensures that subjects have to make choices that are associated with material incentives and is in that sense complementary to vignette designs that are often criticized because they rely on subjects' responses to hypothetical situations. Finally, our design allows studying embeddedness effects on trustfulness as well as trustworthiness, while much of the earlier research only addresses embeddedness effects on trustfulness.

We conclude by returning to our findings that both dyadic and network control through sanction opportunities of trustors affect trustee behavior, while there is evidence for dyadic control effects on trustor behavior but no evidence for network control effects on trustor behavior. These findings correspond to the results from survey research on trust problems in buyer-supplier relations (Buskens 2002: chapter 5; Batenburg et al. 2003; Rooks et al. 2006). This survey research focuses on how embeddedness affects trustfulness of buyers in the sense of investing less in costly contractual safeguards that mitigate bad performance, including opportunistic behavior, of suppliers such as delivery of inferior quality, delivery delays, or bad service. Also, this survey research focuses on how embeddedness affects supplier performance and thus supplier trustworthiness. Results indicate that suppliers react to network control.

Specifically, more sanction opportunities for buyers through contacts and information exchange with other clients of their own suppliers, for example, are associated with better supplier performance (Rooks et al. 2006). This is in line with our experimental result that trustworthiness increases with network control through sanction opportunities of trustors. Also, dyadic control based on buyers' sanction opportunities through expected future business with the supplier affects buyers' investments in costly contractual safeguards (Batenburg et al. 2003). However, there is hardly any empirical evidence for effects of network control on buyer behavior (Buskens 2002). Again, these findings concur with our experimental results on the effects of control on trustor behavior.

Buskens (2002: 152-161) provides arguments and empirical evidence suggesting that the lack of network control effects on buyer behavior is partly due to design, data, and measurement problems of the survey, including possible endogeneity of network embeddedness and sample selectivity. However, these are not plausible arguments for the lack of network control effects on trustor behavior in our experiment. Alternatively, one might wonder whether the experimental findings on network control and trustfulness indicate limits of strategic rationality of trustors. Consider the situation of the trustee (or, respectively, the supplier). He has a good reason to react to dyadic *as well as* network control when he anticipates that his present trustworthiness might affect future trustfulness of the same or other trustors. Similarly, the trustor has a good reason to react to dyadic control in the sense of her *own* future sanction opportunities when she anticipates that the trustee anticipates how his present trustworthiness will affect her *own* future trustfulness. However, the trustor needs to reason "more steps ahead" before having a good reason to react to her network control through future sanction opportunities of *other* trustors. Namely, she has to anticipate that the trustee anticipates how his present trustworthiness will affect future trustfulness of *other* trustors and that *other* trustors will condition their trustfulness on the trustee's present trustworthiness. It may be less likely that actors reason so many steps ahead, especially in unfamiliar settings such as the one in our experiment (see Binmore 1998: chap. 0.4.2; Kreps 1990b: Camerer 2003 for general arguments in this direction). Also, there may be individual variation with respect to the ability to reason ahead. For example, Palacios-Huerta and Volij (2008, 2009) offer empirical evidence from field data and experiments indicating that experienced actors' behavior corresponds more closely to equilibrium behavior than the behavior of less experienced actors, and that experienced actors are better able to reason ahead. They also offer evidence for individual variation in the ability to reason ahead (see Levitt et al. 2010 for a critical discussion).

Future research on embeddedness effects on trust could further explore such conjectures in various ways. For example, if the conjectures are correct, we would expect the effects of network control on trustor behavior to be more easily found when trustors play repeated Trust Games with information exchange between trustors more often and specifically when they are also in the role of the trustee in some of those repeated games (see Bednar and Page 2007 and Bednar et al. 2008 for related arguments and some preliminary empirical evidence in a different context on “spillover effects” between games). The design of our experiment offers some, though limited, empirical evidence on this issue. Namely, each subject played three repeated Trust Games and was a trustor in two of these games and a trustee in the third game. We can thus check whether there are indications for effects of network control on trustfulness for those trustors who had been in the role of the trustee in an earlier game. Analysis of only the cases for trustors who were involved as trustee before indeed reveals a weak main effect of the condition with full information exchange between trustors. Nevertheless, we do not find differences in end-game effects between trustor 1 and trustor 2 according to Hypothesis 6. The available evidence from our experiment on this issue is thus promising but inconclusive. Designs with more possibilities for subjects to gain experience with the setting are clearly needed to further investigate this issue. Such designs are currently being developed and implemented. Also, such designs will benefit from including measurements for subjects’ ability to reason ahead as well as questioning subjects at the end of the experiment on their subjective reasons for choices they made in the Trust Games.

References

Anderhub, V., Engelmann, D., Güth, W., 2002. An experimental study of the repeated trust game with incomplete information. Journal of Economic Behavior and Organization 48, 197-216.

Barrera, D., Buskens, V., 2009. Third-party effects. In: Cook, K.S., Snijders, C., Buskens, V., Cheshire, C. (Eds.) eTrust. Forming Relationships in the Online World. Russell Sage, New York, pp. 37-72.

Batenburg, R.S., Raub, W., Snijders, C., 2003. Contacts and contracts: temporal embeddedness and the contractual behavior of firms. Research in the Sociology of Organizations 20, 135-188.

Bednar, J., Chen, Y., Liu, T.X., Page, S., 2008. Behavioral spillovers in multiple games: an experimental study. Mimeo, University of Michigan.

Bednar, J., Page, S., 2007. Can game(s) theory explain culture? The emergence of cultural behavior within multiple games. Rationality and Society 19, 65-97.

Binmore, K., 1998. Game Theory and the Social Contract, Volume 2: Just Playing. MIT Press, Cambridge, MA.

Blumberg, B.F., 1997. Das Management von Technologiekooperationen: Partnersuche und Verhandlungen mit dem Partner aus empirisch-theoretischer Perspektive. Thesis Publishers, Amsterdam.

Bolton, G.E., Katok, E., Ockenfels, A., 2004. How effective are online reputation mechanisms? An experimental study. Management Science 50, 1587-1602.

Bolton, G.E., Ockenfels, A., 2000. ERC: a theory of equity, reciprocity, and competition. American Economic Review 90, 166-193.

Bolton, G.E., Ockenfels, A., 2009. The limits of trusts in economic transactions: investigations of perfect reputation systems. In: Cook, K.S., Snijders, C., Buskens, V., Cheshire, C. (Eds.) eTrust. Forming Relationships in the Online World. Russell Sage, New York, pp. 15-36.

Bower, A., Garber, S., Watson, J.C., 1997. Learning about a population of agents and the evolution of trust and cooperation. International Journal of Industrial Organization 15, 165-190.

Brandts, J., Figueras, N., 2003. An exploration of reputation formation in experimental games. Journal of Economic Behavior and Organization 50, 89-115.

Brown, G., 1951. Iterative solution of games by fictitious play. In: Koopmans, T.C. (Ed.) Activity Analysis of Production and Allocation. Wiley, New York.

Burt, R.S., Knez, M., 1995. Kinds of third-party effects on trust. Rationality and Society 7, 1-36.

Bush, R. Mosteller, F., 1955. Stochastic Models for Learning. Wiley, New York.

Buskens, V., 2002. Trust and Social Networks. Kluwer, Boston, MA.

Buskens, V., 2003. Trust in triads: effects of exit, control, and learning. Games and Economic Behavior 42, 235-252.

Buskens, V., Raub, W., 2002. Embedded trust: control and learning. Advances in Group Processes 19, 167-202.

Buskens, V., Raub, W., 2010. Rational choice research on social dilemmas. Forthcoming in Wittek, R., Snijders, T.A.B., Nee, V. (Eds.) Handbook of Rational Choice Social Research. Russell Sage, New York.

Buskens, V., Weesie, J., 2000. An experiment on the effects of embeddedness in trust situations: buying a used car. Rationality and Society 12, 227-253.

Camerer, C.F., 2003. Behavioral Game Theory. Experiments in Strategic Interaction. Russell Sage, New York.

Camerer, C.F., Ho, T., 1999. Experience-weighted attraction learning in normal form games. Econometrica 67, 827-874.

Camerer, C.F., Ho, T., Chong, J.-K., 2002. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. Journal of Economic Theory 104, 137-188.

Camerer, C.F., Weigelt, K., 1988. Experimental tests of a sequential equilibrium reputation model. Econometrica 56, 1-36.

Cheung, Y.-W., Friedman, D., 1997. Individual learning in normal form games: some laboratory results. Games and Economic Behavior 19, 46-76.

Coleman, J.S., 1990. Foundations of Social Theory. Belknap Press of Harvard University Press, Cambridge, MA.

Cournot, A., 1838 (1960). Recherches sur les Principes Mathématiques de la Théorie des Richesses. Translated into English by N. Bacon as Researches in the Mathematical Principles of the Theory of Wealth. Haffner, London.

Dasgupta, P., 1988. Trust as a commodity. In: Gambetta, D. (Ed.) Trust: Making and Breaking Cooperative Relations. Blackwell, Oxford, pp. 49-72.

Erev, I., Roth, A.E., 1998. Predicting how people play games: reinforcement learning in experimental games with unique, mixed-strategy equilibria. American Economic Review 88, 848-881.

Falk, A., Heckman, J.J., 2009. Lab experiments are a major source of knowledge in the social sciences. Science 326, 535-538.

Fehr, E., Schmidt, K.M., 1999. A theory of fairness, competition, and cooperation. Quarterly Journal of Economics 114, 817-868.

Fischbacher, U., 2007. z-Tree - Zurich Toolbox for Readymade Economic Experiments – Experimenter's Manual. Experimental Economics 10, 171-178.

Fudenberg, D., Levine, D.K., 1998. The Theory of Learning in Games. MIT Press, Cambridge, MA.

Gautschi, T., 2000. History effects in social dilemma situations. Rationality and Society 12, 131-162.

Goyal, S., 2007. Connections. An Introduction to the Economics of Networks. Princeton University Press, Princeton, NJ.

Granovetter, M., 1985. Economic action and social structure: the problem of embeddedness. American Journal of Sociology 91, 481-510.

Greiner, B., 2004. The Online Recruitment System ORSEE 2.0. A Guide for the Organization of Experiments in Economics. University of Cologne, Working Paper Series in Economics 10.

Harsanyi, J.C. 1967/68. Games with incomplete information played by 'Bayesian' players I-III. Management Science 14, 159-182, 320-334, 486-502.

Huber, P., 1967. The behavior of the Maximum Likelihood estimates under nonstandard conditions. Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability 1, 221-233.

Jackson, M.O., 2008. Social and Economic Networks. Princeton University Press, Princeton, NJ.

Kreps, D.M., 1990a. Corporate culture and economic theory. In: Alt, J.E., Shepsle, K.A. (Eds.) Perspectives on Positive Political Economy. Cambridge University Press, Cambridge, pp. 90-143.

Kreps, D.M., 1990b. Game Theory and Economic Modelling. Clarendon Press, Oxford.

Kreps, D.M., Wilson, R., 1982. Sequential equilibria. Econometrica 50, 863-894.

Levitt, S.D., List, J.A., Reiley, D.H., 2010. What happens in the field stays in the field: professionals do not play minimax in laboratory experiments. Forthcoming in Econometrica.

Lorenz, E.H. 1988. Neither friends nor strangers: informal networks of subcontracting in French industry. In: Gambetta, D. (Ed.) Trust: Making and Breaking Cooperative Relations. Blackwell, Oxford, pp. 94-107.

Macy, M.W., Flache, A., 2002. Learning dynamics in social dilemmas. Proceedings of the National Academy of Sciences U.S.A. 99, 7229-7236.

Neral, J., Ochs, J., 1992. The sequential equilibrium theory of reputation building: a further test. Econometrica 60, 1151-1169.

Palacios-Huerta, I., Volij, O., 2008. Experientia docet: professionals play minimax in laboratory experiments. Econometrica 76, 71-115.

Palacios-Huerta, I., Volij, O., 2009. Field centipedes. American Economic Review 99, 1619-1935.

Rasmusen, E., 2007. Games and Information: An Introduction to Game Theory. 4th ed., Blackwell, Oxford.

Raub, W., 1988. An analysis of the finitely repeated Prisoner's Dilemma. European Journal of Political Economy 5, 367-380.

Raub, W., Weesie, J., 1990. Reputation and efficiency in social interactions: an example of network effects. American Journal of Sociology 96, 626-654.

Rogers, W.M. 1993. Regression standard errors in clustered samples. Stata Technical Bulletin 13, 19-23.

Rooks, G., Raub, W., Selten, R., Tazelaar, F., 2000. Cooperation between buyer and supplier: effects of social embeddedness on negotiation effort. Acta Sociologica 43, 123-137.

Rooks, G., Raub, W., Tazelaar, F., 2006. Ex post problems in buyer-supplier transactions: effects of transaction characteristics, social embeddedness, and contractual governance. Journal of Management and Governance 10, 239-276.

Roth, A.E., 1995. Introduction to experimental economics. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, NJ, pp. 3-109.

Roth, A.E, Erev, I., 1995. Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior 8, 164-212.

Snijders, C., 1996. Trust and Commitments. Thesis Publishers, Amsterdam.

Snijders, C., Keren, G., 2001. Do you trust? Whom do you trust? When do you trust? Advances in Group Processes 18, 129-160.

Snijders, T.A.B., Bosker, R.J., 1999. Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling. Sage, Thousand Oaks, CA.

Williamson, O.E., 1996. The Mechanisms of Governance. Oxford University Press, New York.

Yamagishi, T., Yamagishi, M., 1994. Trust and commitment in the United States and Japan. Motivation and Emotion 18, 129-166.

Table 1. Trustfulness and trustworthiness, per round, and per information condition.

Round		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	**Total**[1]
Total																	
Trustfulness	N	133	124	123	117	118	120	122	117	116	121	116	114	108	89	37	1675
	%[2]	92	86	85	81	82	83	85	81	81	84	81	79	75	62	26	78
Trustworthiness	N	117	106	104	95	98	103	100	100	108	100	93	96	86	54	14	1374
	%[3]	88	85	84	81	83	86	82	85	93	83	80	84	80	61	38	82
Full information exchange between trustors																	
Trustfulness	N	70	66	64	63	67	62	63	64	64	66	64	65	63	52	20	913
	%[2]	97	92	89	88	93	86	88	89	89	92	89	90	88	72	28	85
Trustworthiness	N	65	59	60	57	59	54	57	59	63	60	57	60	56	39	9	814
	%[3]	93	89	94	90	88	87	90	92	98	91	89	92	89	75	45	89
No information exchange between trustors																	
Trustfulness	N	63	58	59	54	51	58	59	53	52	55	52	49	45	37	17	762
	%[2]	88	81	82	75	71	81	82	74	72	76	72	68	63	51	24	71
Trustworthiness	N	52	47	44	38	39	49	43	41	45	40	36	36	30	15	5	560
	%[3]	83	81	75	70	77	84	73	77	87	73	69	73	67	41	29	73

[1] Total number of games played = 2160

[2] Percentage of games in which the trustor is trustful, as percentage of all games

[3] Percentage of games in which the trustee is trustworthy, as percentage of games in which the trustor is trustful

Table 2. Three-level logistic regression of TRUSTFULNESS (2160 decisions by 144 trustors in 72 triads).

		Baseline model		Full model	
	Hyp.	Coeff.	Std. err.	Coeff.	Std. err.
FULL INFORMATION	+	0.983**	0.332	-0.531	0.652
REPEATEDGAME2		0.381	0.379	-0.015	0.258
REPEATEDGAME3		0.911*	0.385	0.491	0.290
TRUSTOR2		-0.059	0.124	0.108	0.204
PASTHONOR_OWN	+			1.781**	0.179
PASTABUSE_OWN	−			-1.600**	0.242
PASTHONOR_OTHER	+			0.647*	0.253
PASTABUSE_OTHER	−			-1.266**	0.404
FUTURE	+			0.077**	0.029
FUTUREFULL	−			0.060	0.052
ROUND14	−			-0.444	0.496
ROUND15	−			-1.302*	0.539
ROUND14FULL				-0.708	0.833
ROUND15FULL				-2.331**	0.800
ROUND14TR2				-0.744	0.668
ROUND15TR2				-1.434	0.766
ROUND14TR2FULL	−			0.320	1.041
ROUND15TR2FULL	−			1.528	1.030
ROUND1REPEATEDGAME1				1.341*	0.563
ROUND1REPEATEDGAME2				1.563*	0.619
ROUND1REPEATEDGAME3				1.840*	0.815
CONSTANT		0.784*	0.321	-0.249	0.389
Variance triad level		1.364	0.332	0.172	0.165
Variance trustor level		0.251	0.114	0.331	0.202
Variance decision level		3.290		3.290	
Loglikelihood		-972.510		-615.191	

*, ** indicate significance at $p < 0.05$ and $p < 0.01$, respectively (two-sided tests).

Table 3. Three-level logistic regression of TRUSTWORTHINESS (1542 decisions with 144 trustors by 72 trustees in 72 triads).

		Baseline model		Full model	
	Hyp.	Coeff.	Std. err.	Coeff.	Std. err.
FULL INFORMATION	+	1.405**	0.317	2.263**	0.577
REPEATEDGAME2		0.295	0.385	0.312	0.512
REPEATEDGAME3		0.439	0.383	0.378	0.509
TRUSTOR2		-0.513	0.147	0.007	0.169
FUTURE	+			0.080*	0.031
FUTUREFULL	−			-0.073	0.050
ROUND14	−			-1.735**	0.603
ROUND15	−			-2.889**	0.868
ROUND14FULL				1.115	0.902
ROUND15FULL				-0.071	1.201
ROUND14TR2				-0.372	0.798
ROUND15TR2				0.825	1.299
ROUND14TR2FULL	−			-1.431	1.113
ROUND15TR2FULL	−			-3.596*	1.770
CONSTANT		0.681*	0.318	0.307	0.482
Variance triad level		1.289	0.352	2.492	0.663
Variance trustor level		0	0	0	0
Variance decision level		3.290		3.290	
Loglikelihood		-675.315		-593.339	

*, ** indicate significance at $p < 0.05$ and $p < 0.01$, respectively (two-sided tests).

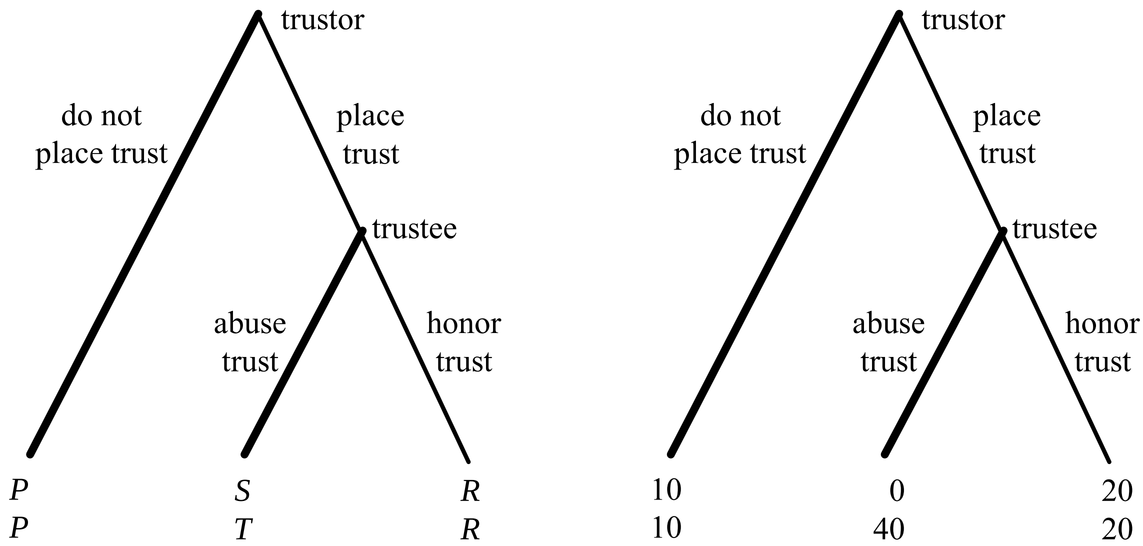

Figure 1. Extensive form of a Trust Game. $T > R > P > S$. The right-hand Trust Game is the numerical example used in the experiment. Bold lines indicate the equilibrium path of play.

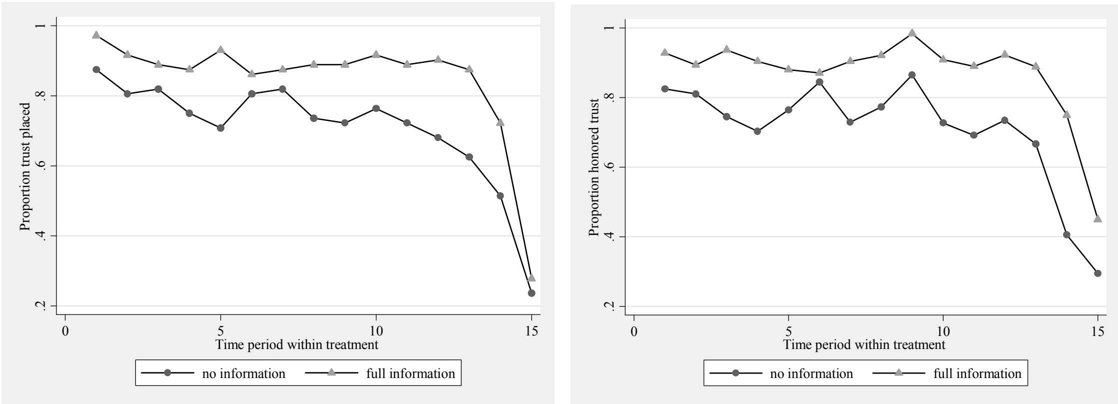

Figure 2. Proportion of games in which trust was placed (left) and trust was honored (right), per round, and per experimental condition.