Trust in triads

Buskens, Vincent

Published in:
Games and Economic Behavior

DOI:
10.1016/S0899-8256(02)00563-8

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Games and Economic Behavior 42 (2003) 235–252

GAMES and Economic Behavior

www.elsevier.com/locate/geb

Trust in triads: effects of exit, control, and learning

Vincent Buskens

Department of Sociology, Utrecht University, The Netherlands

Received 11 January 2000

Abstract

This paper provides theoretical background for some effects of social networks on trust. We study the implications of a model with rational actors in two settings with three actors. In the first setting, there are two trustees who are involved in transactions with one truster implying that the truster has an *exit* option. In the second setting, two trusters play with one trustee, which gives the trusters options for *voice*, i.e., complaining and informing each other about the trustee's behavior. We compare these models with a baseline model in which there is only one truster and one trustee. It turns out that the opportunities for placing and honoring trust do not change for the exit model compared to the baseline model. The opportunities for trust in the voice model differ from the baseline model only if both trusters inform each other at a rate that is high enough. Only if the possibilities for receiving information *and* transmitting information are large enough for both trusters, trust will increase due to the information exchange possibilities in the voice model.

JEL classification: C72

Keywords: Trust; Social networks; Non-cooperative game theory

1. Introduction

The relation between social networks and trust is a quite complex one, because social networks may have different types of effects on trust. For example, networks can provide options for exit out of a relation (Lahno, 1995), for obtaining information from previous behavior of other actors (learning, see Buskens, 2002), or for controlling partners through reputational sanctions when they act untrustworthy (Coleman, 1990; Kreps, 1990; Raub and Weesie, 1990; Buskens, 2002). In this paper, I model these three effects in

E-mail address: v.buskens@fss.uu.nl.

0899-8256/03/$ – see front matter
doi:10.1016/S0899-8256(02)00563-8

similar models that are all based on a well-known game-theoretic model with incomplete information to analyze the finitely repeated Prisoner's Dilemma (see Kreps et al., 1982).

In a trust situation, a truster has to decide first whether or not to trust a trustee. Placing trust allows the trustee to choose between honoring and abusing trust, which would not have been possible if the truster would not have placed trust. The truster regrets placing trust if trust is abused, but benefits from honored trust. The trustee can earn an extra profit from abusing trust in a transaction. Therefore, if a transaction is happening in isolation, i.e., with a trustee that is unknown to the truster, the truster and trustee do not expect to meet at any time after the transaction, and the trustee is expected to take this extra profit. Consequently, the truster will not place trust in such a situation. Formally, a trust situation can be represented by a Trust Game (Dasgupta, 1988; Kreps, 1990), which will be introduced in Section 2.

I am convinced that many exchange relations or transactions resemble a trust situation as given above. For example, an actor who wants to buy a used car knows that the dealer has an incentive to sell the car for a price that is too high. The dealer might conceal essential information about the history of the car, for example, whether the car has been involved in a major accident that has caused vital damage to the car. The buyer is assumed to be unable to deduce this information by inspection of the car. However, the buyer is also uncertain about the extent to which the dealer has an incentive to sell at a high price. The dealer might be concerned about future business with this buyer or acquaintances of this buyer. Moreover, the dealer might just feel guilty if he would conceal information. Consequently, the buyer is not only uncertain about the quality of the car (which creates the trust problem), but also about the precise incentives of the dealer.

Exchanges or transactions among actors hardly ever happen in isolation. Most transactions are embedded in a larger social setting, for example, because actors have more transactions with each other (*temporal embeddedness*) or because third parties are connected to the actors in a transaction (*network embeddedness*). These two types of embeddedness might affect the behavior of the actors involved in a transaction (Raub and Weesie, 2000). In this paper, I want to concentrate on three effects of third parties using the smallest possible networks that exceed the dyadic level: *triads*. This provides the opportunity to reach some analytic results, which are difficult to obtain for larger systems. Moreover, it provides possibilities for testing the theory in laboratory experiments.

First, I discuss a baseline model in which only one truster and one trustee are involved in a finite number of transactions. Second, the model will be extended with an *exit* option for trusters (Hirschman, 1970, Chapter 2). An exit option increases the set of feasible sanctions for the truster because a truster cannot only sanction the trustee by withholding trust, but she can switch to another trustee as well. Third, a *voice* option is incorporated for trusters (Hirschman, 1970, Chapter 3). In this case, there are two trusters who are involved in transactions with the trustee and they can communicate about the behavior of the trustee. This provides the truster with additional opportunities to *control* the trustee, because she can inform the other truster about the behavior of the trustee, and the second truster may refrain from placing trust as a result of this information. Moreover, the trusters can *learn* about the trustee's incentives to abuse trust from each others experiences with the trustee in the past. Because the learning opportunities of one truster coincide with the control opportunities of the other truster, I can distinguish between these two kinds of effects only by allowing asymmetric information flows between the two trusters. In this way, I hope to disentangle whether trust can be facilitated better by control or learning and what the relative impact of these two effects is if they are combined.

The paper is outlined as follows. Section 2 summarizes the theory on the finitely repeated Trust Game with incomplete information, which is called the baseline model and is discussed as a reference to be compared with the other models.[1] In Section 3, an exit option is introduced. Section 4 analyzes the finitely repeated Trust Game with two trusters who might communicate about the behavior of the trustee between the periods of play. Section 5 summarizes the main findings and testable hypotheses that follow from the models presented and gives indications for further theoretical developments. Finally, some comments are made about possibilities to test these models.

2. The baseline model

The model developed here closely resembles reputation models in the economic literature, in particular the models developed by Kreps and Wilson (1982a) on the "chain-store paradox" and by Kreps et al. (1982) on the finitely repeated Prisoner's Dilemma. The constituent game here is the Trust Game as shown in Fig. 1. The Trust Game is a one-sided version of the Prisoner's Dilemma. I assume that there are two types of trustees: "friendly" and "payoff-maximizing." Both types are utility maximizers, but friendly trustees will always feel guilty to such an extent that $u_2(P_2) < u_2(T_2) < u_2(R_2)$ (cf. Güth and Kliemt, 1994; Snijders, 1996).[2] Consequently, friendly trustees will never abuse trust. Since friendly trustees have no short-term incentive to abuse trust, they certainly will not have a long-term incentive to abuse trust. One could model these trustees as if they do not have the option to abuse trust, which would essentially lead to the same equilibria. Payoffs represent

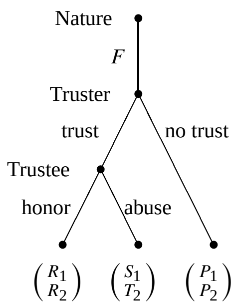

Fig. 1. Extensive form of a Trust Game with incomplete information, where $R_i > P_i$ $(i = 1, 2)$, $P_1 > S_1$, $T_2 > R_2$, and F is the distribution over the types of trustees.

[1] This model is discussed in detail by Bower et al. (1997) and is mentioned earlier in a more informal way by Camerer and Weigelt (1988), Dasgupta (1988), and Neral and Ochs (1992).

[2] It is known from earlier experiments that there is a considerable number of trustees who actually honor trust also if the Trust Game or similar games are played only once (McKelvey and Palfrey, 1992; Berg et al., 1995; Snijders, 1996; Güth et al., 1997; Snijders and Keren, 1999).

utility for the payoff-maximizing trustee. I assume that also the trusters' payoffs represent utility. Therefore, with some abuse of notation, I will denote the utilities for the trusters and payoff-maximizing trustees by the payoffs as given in the Trust Game.[3]

The game starts with a move by Nature deciding which type of trustee is going to play. The trustee is "friendly" with a probability π and "payoff-maximizing" with a probability $1 - \pi$. These proportions are common knowledge. The trustee knows his type, but the truster does not observe the type of the trustee. No discounting is assumed for payoffs received in later periods of the game. The periods are labeled backwards, i.e., the last period is period 1 and the first period is period N. Furthermore, π_n is the belief of the truster that the trustee is of the friendly type at the start of period n; p_n is the probability that a truster places trust in period n; q_n is the probability that a payoff-maximizing trustee honors trust in period n. I define:

$$\mathrm{RISK} = \frac{P_1 - S_1}{R_1 - S_1} \quad \text{and} \quad \mathrm{TEMP} = \frac{T_2 - R_2}{T_2 - P_2}. \tag{1}$$

RISK represents the risk for trusters to place trust and TEMP represents the temptation for the trustee to abuse trust (see Snijders, 1996; Snijders and Keren, 1999). The following beliefs and strategies form a sequential equilibrium (Kreps and Wilson, 1982b) for the finitely repeated Trust Game with incomplete information as described above.

Beliefs of the truster

- If the truster does not place trust in period $n + 1$, then $\pi_n = \pi_{n+1}$.
- If the truster places trust in period $n + 1$ and the trustee honors trust in that period, the truster updates beliefs: $\pi_n = \max(\mathrm{RISK}^n, \pi_{n+1})$.
- If the truster places trust in period $n + 1$ and the trustee abuses trust, $\pi_n = 0$.

Strategy of the truster

If $\pi_n > \mathrm{RISK}^n$, the truster places trust in period n. If equality holds, the truster randomizes with a probability $p_n = \mathrm{TEMP}$. Otherwise, the truster does not place trust.

Strategy of a payoff-maximizing trustee

- If $\pi_n \geqslant \mathrm{RISK}^{n-1}$, a payoff-maximizing trustee honors trust.
- If $\pi_n < \mathrm{RISK}^{n-1}$, a payoff-maximizing trustee honors trust with a probability $q_n = (\pi_n/(1 - \pi_n))(1/\mathrm{RISK}^{n-1} - 1)$.

Theorem 2.1. *The strategies and beliefs above constitute a sequential equilibrium in the finitely repeated Trust Game with friendly and payoff-maximizing trustees.*

[3] It is assumed that the payoffs are the same in all the periods. The analysis can be generalized for arbitrary payoffs in all periods, but this will considerably complicate the notation.

Proof. The proof for $N = 2$ is given by Bower et al. (1997). They also provide the result for $N > 2$, which follows from an induction argument. □

This equilibrium consists of three phases. The game starts with a number of periods in which trust is placed and honored. At a certain period, the trustee starts to randomize. In this second phase, both truster and trustee randomize until the truster does not place trust or the trustee abuses trust. After any instance of no trust or abuse of trust the third and last phase starts in which there will be no more trust until the end of the game.

The equilibrium demonstrates that whether or not a truster trusts the trustee depends mainly on RISK, the number of periods to be played until the end of the game, and the proportion of friendly trustees in the total population. The higher the risk of placing trust for the truster, the smaller the number of periods to be played, and the smaller the ex-ante probability that a trustee is friendly, the shorter the first (trust) phase of the game will be. When the payoff-maximizing trustee starts to randomize, the truster learns in the sense that she updates her belief about the probability that she is playing with a friendly trustee. This probability increases gradually as long as trust is honored and becomes zero as soon as trust is abused.

It is striking that the payoffs of the truster incorporated in RISK are of major importance to determine how close the game can approach the end before trust breaks down. This contrasts with well-known results for the infinitely repeated Trust Games with complete information and discounting. In the latter case, trust is explained completely by the discount factor δ of the trustee and the payoffs of the trustee. Namely, δ should be larger than TEMP (see, for example, Kreps, 1990). TEMP plays only a role in the randomization of the trusters in the game analyzed here. If TEMP is larger, the probability that the truster places trust in the randomization phase is larger. This is in a sense counterintuitive, since this implies that trusters place more trust if the temptation of the trustee for abusing trust is larger. The reason is that the truster's randomization probability is chosen such that the trustee is indifferent in the period before. Therefore, the larger the temptation for the trustee, the higher the probability needs to be that trust is placed in the following period again, to make him indifferent between honoring and abusing trust. In the following two sections, I will discuss changes in the predictions if the truster has an exit option or if there are two trusters who can communicate about the trustee's behavior.

3. The exit model

The first extension is an exit option. Now, one truster plays a finitely repeated Trust Game with two (or more) trustees. Between every two periods of the game, the truster has an additional choice, namely, which trustee she wants to encounter in the next period. It is assumed that a trustee who is not playing with the truster in a period receives a payoff P_2, the non-playing payoff that corresponds with the no-trust payoff. Again the ex-ante probability that a trustee is friendly is π and that a trustee is payoff-maximizing is $1 - \pi$. The costs of switching trustees for the truster are small compared to the payoffs in the game, such that switching costs only affect the truster's behavior if she would be indifferent in case there would be no costs of switching.[4] The truster starts with choosing the trustee she wants to play with and, thereafter, she chooses whether she does or does not trust this trustee.

The main equilibrium for this game can be derived easily by reconsidering the equilibrium for the baseline model. It is known that payoff-maximizing trustees mimic friendly trustees as long as that is better for them. At a certain moment, the payoff-maximizing trustee switches to randomization to convince the truster that he is really friendly. If this randomization leads to an abuse of trust, the truster probably would switch to another trustee, but at that moment it is too late for the truster to experiment with a new trustee because the future is too short, and payoff-maximizing trustees will not honor trust anymore. Formally, since the trustee abused trust in the foregoing period, it holds in the present period n that $\pi < \mathrm{RISK}^n$, which implies that the truster will not trust a trustee she did not encounter before. Thus, as long as all trustees always honor trust, the truster has no reason to change trustees. As soon as she discovers that she is playing with a payoff-maximizing trustee, it is too late to trust any trustee. One might expect that if the trustee's non-playing payoff is smaller than P_2, the sanction from changing trustees increases for the untrustworthy trustee. However, such a payoff change for the trustee does only affect the randomization probability of the truster. Moreover, the probability that the truster places trust after trust is honored in the randomization phase even decreases, because the expected payoff for the trustee after honoring trust should be the same as his expected payoff after abusing trust. Thus, since the expected payoff after abusing trust decreases as a result of the additional sanction, this implies that the expected payoff after honoring trust should decrease as well. The truster will not be better off as a result of this change, because the expected payoff for her in the randomization phase remains P_1. Consequently, the equilibrium strategy for the truster is choosing one of the trustees to begin with and play the equilibrium strategy of the baseline model with this trustee. The chosen trustee will also follow the equilibrium strategy of the baseline model.

It might be clear that this outcome for the inclusion of an exit option is not a satisfying outcome. An exit option increases the sanction opportunities for the truster and decreases the dependence of the truster on the trustee. It should be noted that also in the repeated Trust Game with complete information, an exit option would not have had an effect on the solution. This is often attributed to the fact that all the trustees are the same and there is complete information about their characteristics. I have shown here that assuming incomplete information does not automatically solve this problem. Still, there are many other ways to model exit opportunities, although this easily leads to rather complicated models. One option would be to model a repeated Trust Game with monitoring problems in which the trustee may abuse trust unintentionally, while different trustees have different capabilities, i.e., some trustees abuse trust unintentionally with a higher probability than others. Other authors have developed models for exit in different settings related to

[4] If there would be no costs of switching trustees for the truster, the truster would be indifferent between switching or staying with the same trustee in the early periods of the game (because she did not learn anything yet) or after trust has been abused (because she will not trust any truster). This would cause uncertainty on the side of the trustees about whether the truster will continue to play with him. By imposing some costs of switching, which seems a reasonable assumption, I circumvent this problem.

trust and other cooperation problems (see Schüßler, 1989; Vanberg and Congleton, 1992; Lahno, 1995; Weesie, 1996; Macy and Skvoretz, 1998).

4. The voice model

Now, the baseline model is extended with a second truster. Both trusters play N Trust Games with the trustee. One truster starts and in the following period the other truster plays. Consequently, there will be $2N$ periods of play and the game starts with period $2N$. Between two periods, it is decided by a probabilistic mechanism whether the former truster can transmit information to the other truster about former play. For now assume that, if possible, the former truster informs the latter truster truthfully about what happened in the last period. Later, I will discuss this assumption in some more detail. I will also show that it does not matter for the equilibrium discussed here whether or not the trustee observes communication between the trusters as long as he knows the probabilities for information transmission. The probabilities for information transmission from truster i to truster j are τ_{ij}, $i, j = 1, 2$, $i \neq j$. The transmission probabilities do not need to be the same, so it does not need to be the case that $\tau_{ij} = \tau_{ji}$. This results in a small, asymmetric communication network. The communication probability τ_{12} is called the *outdegree* of truster 1 or the *indegree* of truster 2. Similarly, τ_{21} is the indegree of truster 1 and the outdegree of truster 2. Truster 1 controls the trustee through her outdegree τ_{12}, while she learns about the trustee through her indegree τ_{21}. If trust would be based primarily on what trusters hear about the trustee, truster 1's trust would be affected more by her indegree τ_{21}, while truster 2's trust would be more affected by τ_{12}. If trust would be based more on the potential sanctions imposed on the trustee after abusing trust, outdegree should be more important.

Now, I describe the beliefs of the trusters and, thereafter, the strategies of the players. I will not use different indices for the two trusters. The indices refer to the period in the game. Because truster 1 starts, even indices are related to truster 1, and odd indices are related to truster 2. Two cases need to be distinguished:

Case 1. $\tau_{12} \geqslant \text{TEMP}$ and $\tau_{21} \geqslant \text{TEMP}$;
Case 2. $\tau_{12} < \text{TEMP}$.

For the cases in which $\tau_{12} \geqslant \text{TEMP}$ and $\tau_{21} < \text{TEMP}$, the equilibrium resembles the equilibrium of Case 2, and the qualitative implications are the same as for Case 2.[5]

Beliefs of the trusters

- If a truster does not place trust, does not obtain information from the other truster, or is informed that the other truster did not place trust, beliefs about the trustee do not change.

[5] Complete proofs for all cases can be obtained from the author.

- If a truster knows about any abuse of trust by the trustee, $\pi = 0$ for all subsequent periods of this truster (the truster knows that the trustee is not friendly).
- In any case, if a truster receives information from the other truster about behavior of the trustee, she updates her belief about the probability that the trustee is a friendly trustee to the same value as the belief of the truster who transmits the information to her.

Case 1. If the truster places trust in period $n+1$ and the trustee honors trust in that period, the truster updates her belief to the value $\pi_n = \max(\text{RISK}^n, \pi_{n+1})$.

Case 2. If truster 1 places trust and the trustee honors trust in period $n+1$, she updates her belief such that $\pi_n = \max(\text{RISK}^{(n-1)/2}, \pi_{n+1})$. If truster 1 did not place trust in period $n+1$, but she received information about honored trust in period n from truster 2, she will update her belief to $\pi_{n-1} = \max(\text{RISK}^{(n-1)/2}, \pi_n)$. If truster 2 received information from truster 1 before her period, her belief will not change after her own period. If she did not receive information from truster 1, she updates her belief after honored trust in period $n+1$ to $\pi_n = \max(\text{RISK}^{n/2}, \pi_{n+1})$.

Strategies of the trusters

Case 1. If $\pi_n > \text{RISK}^n$, the truster places trust in period n. If $\pi_n = \text{RISK}^n$, trusters 1 and 2 place trust with a probability TEMP/τ_{21} and TEMP/τ_{12}, respectively. Otherwise, the trusters do not place trust.

Case 2. If $\pi_n > \text{RISK}^{\lceil n/2\rceil}$, the truster places trust in period n. If $\pi_n = \text{RISK}^{\lceil n/2\rceil}$, truster 1 places trust with a probability $(\text{TEMP}(1-\tau_{12}) - \tau_{12}\text{TEMP} + \tau_{12}^2)/((1-\tau_{12}) - \tau_{12}\text{TEMP} + \tau_{12}^2)$ if she placed trust in her previous period, but she will not place trust if she did not place trust in her previous period. If $\pi_n = \text{RISK}^{\lceil n/2\rceil}$, truster 2 places trust if she just obtained information that truster 1 did not place trust and she places trust with a probability $(\text{TEMP} - \tau_{12})/(1-\tau_{12})$ if she did not receive any information from truster 1. Otherwise, the trusters do not place trust.

Strategy of a payoff-maximizing trustee

Case 1. If $\pi_n < \text{RISK}^{n-1}$, the trustee honors trust with a probability $q_n = (\pi_n/(1-\pi_n))(1/\text{RISK}^{n-1} - 1)$. If $\pi_n \geqslant \text{RISK}^{n-1}$, a payoff-maximizing trustee honors trust.

Case 2. If $\pi_n \geqslant \text{RISK}^{(n-2)/2}$, the trustee honors trust placed by truster 1. If $\pi_n < \text{RISK}^{(n-2)/2}$, the trustee honors trust placed by truster 1 with probability $q_n = (\pi_n/(1-\pi_n))(1/\text{RISK}^{(n-2)/2} - 1)$. The trustee repeats his move from period n with truster 1 in period $n-1$ with truster 2. If truster 1 did not place trust in period n, the trustee plays the move he would have played in period n in period $n-1$ with truster 2.

Theorem 4.1. *Considering the beliefs and strategies described above*:

- *The strategies and beliefs are a sequential equilibrium in the finitely repeated Trust Game with two trusters.*
- *In Case* 1*, if* $\pi < \mathrm{RISK}^{2N-1}$*, there is one other sequential equilibrium, namely, never placing trust by both trusters and always abusing trust by the trustee.*
- *Otherwise, if* $\pi \neq \mathrm{RISK}^n$ *for* $n \leqslant 2N$*, every sequential equilibrium for Case* 1 *has on-the-equilibrium-path strategies as described previously.*
- *For Case* 2*, there exists no equilibrium for which the trustee starts randomizing more than one period later than in the equilibrium described here.*

Proof. The proof of this theorem is presented in Appendix A. □

The most important substantive finding of the last theorem is that trust increases with the communication opportunities if and only if both trusters transmit information at a high rate, i.e., if $\tau_{12} \geqslant$ TEMP and $\tau_{21} \geqslant$ TEMP (Case 1). As in the baseline model, the equilibrium consists for both cases of three phases. If $\tau_{12} = \tau_{21} = 1$, the equilibrium exactly corresponds with the equilibrium in the game with one truster. In the first phase, trust is placed and honored. After some time, the trustee starts to randomize. Thereafter, the actors remain in the randomization phase only if trust is placed, trust is honored, *and* information is communicated between the trusters. The equilibrium is visualized in Fig. 2. Table 1 provides the explanation for this figure as well as for Fig. 3.[6]

The equilibrium implies that the trustee remains trustworthy with as many periods left as in the baseline model. Consequently, both trusters can trust until they have each only half of the periods left compared to the baseline model. This implies that, compared to the game in which both trusters play in isolation with the trustee, there is considerably more trust in

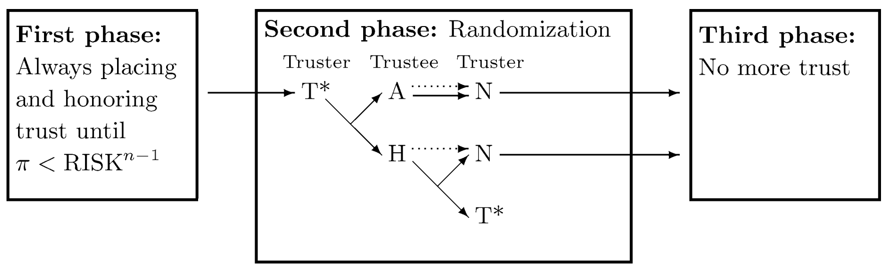

Fig. 2. Equilibrium play in the finitely repeated Trust Game with two trusters and one trustee (Case 1; note that the two trusters are completely interchangeable in this case; further explanation can be found in Table 1 and the main text).

[6] There is another equilibrium in one instance because truster 2 has to start randomizing in her first move. Her expected payoff is the same if she never places trust. If truster 2 does not randomize, truster 1 is better off not placing trust in her first move as well. This equilibrium is weakly Pareto inferior to the first one, because both truster 1 and the trustee are worse off. Uniqueness in all other situations follows from the fact that the game is similar to the baseline model and that both trusters need the randomization phase to sustain trust in the earlier periods of the game.

Table 1
Legend belonging to the figures

- T indicates that the truster places trust; N that the truster does not place trust; H indicates that the trustee honors trust; and A that the trustee abuses trust.
- Normal arrows indicate that the actor who is going to play knows what happened in the last move. Dotted arrows indicate that he or she does not know what happened in the last move. Of course, the trustee always knows what the truster did in her last move. In the baseline model, the truster is always informed about the last move of the trustee as well. Information is not an issue for arrows that enter or leave a boxes.
- Arrows that split in two parts indicate randomization moves.
- In all positions within a figure indicated by T*, the game continues in the same way. This is also true for positions indicated with $N^{\dagger}$.

this game. Again, the higher RISK and the smaller π, the earlier the randomization phase will start. The probability that a truster places trust in the randomization phase increases with TEMP and decreases with the probability that she receives information from the other truster (τ_{12}, τ_{21}). Although these outcomes might be counterintuitive, they are in line with the results for the baseline model. Since the randomization probabilities of the trusters are chosen such that the trustee is indifferent, the truster has to choose higher probabilities of placing trust if the situation for the trustee is better, i.e., if TEMP is larger and if communication between the trusters is less frequent. The net result is that the probabilities to go from the randomization phase to the "no-trust" phase are exactly the same as in the baseline model.

One can argue now that truthful information transmission is beneficial for both trusters in this case, because the equilibrium is built on truthful information exchange among the trusters. The trustee would stop honoring trust earlier if he cannot be sure that the trusters will exchange information truthfully, especially in the randomization phase. The trustee does not need to observe actual communication between the trusters, because trusters do not place trust anymore as soon as information is not transmitted in the randomization phase. Therefore, the trustee can infer whether or not information has been transferred from the actions of the trusters. The trusters only need to communicate what happened in the last period, because what happen in earlier periods can be derived from the last period or the information does not influence the behavior of the truster. The last two observations are also true for Case 2.

In Case 2, truster 1 transmits information at a low rate. It does not matter whether truster 2 transmits information at a high or a low rate. Now, the equilibrium resembles the equilibrium of the baseline model. The randomization phase starts for both trusters with the same number of periods left as in the baseline model. If $\tau_{12} = 0$, the equilibrium reduces to the equilibrium of the baseline model for each of the trusters, whatever the value of τ_{21} is. The strategy of the trustee prescribes that he starts randomizing in a period with truster 1 and he repeats the move that is the outcome of this randomization in the following period with truster 2 (see Fig. 3). By starting randomizing in a period with truster 1, he has a relatively large probability to abuse trust twice, and obtain the T_2 payoff twice. As before it can be seen that trust can be placed longer if RISK is smaller. The randomization probabilities of the trusters are smaller than or equal to TEMP, and the extent to which they have to be smaller depends on τ_{12}. Again, the randomization probabilities of the

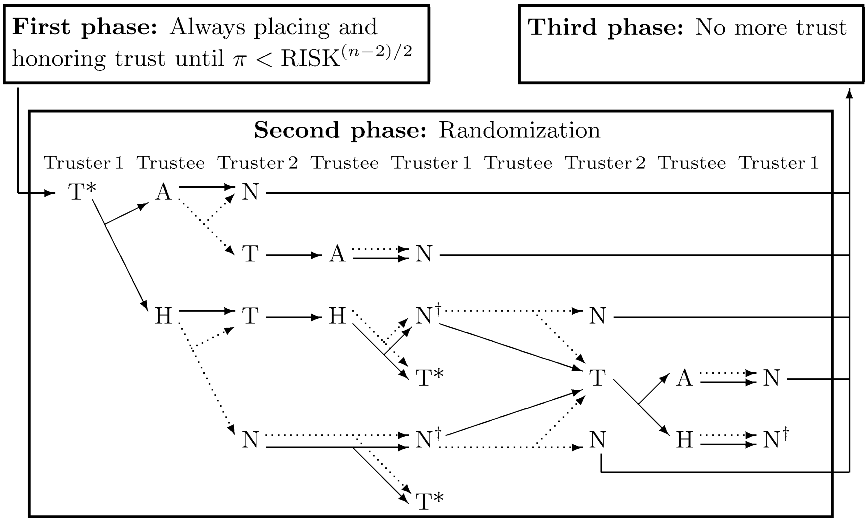

Fig. 3. Equilibrium play in the finitely repeated Trust Game with two trusters and one trustee (Case 2; further explanation can be found in Table 1 and the main text).

trusters increase with TEMP. The randomization probability of truster 2 decrease with τ_{12}, while the randomization probability of truster 1 is u-shaped in τ_{12}. The randomization probabilities do not depend on τ_{21}.

In Case 2, information transmission from truster 2 to truster 1 is worthless. Due to the strategy of the trustee, the information of truster 2 is not new for truster 1 or she will not use the information to adapt her behavior. Figure 3 demonstrates that after trust of truster 1 is honored, the behavior of truster 1 in her next period does not depend on whether or not truster 2 placed trust neither on whether or not information is communicated between the trusters. Truster 2 might profit from the information she receives, because she does not need to randomize if she obtains information about honored trust, and she certainly will receive R_1 in these periods. Truster 2 can randomize or trust again even after she did not place trust before, if she receives information that the trustee has honored trust placed by the other truster in the last period. Finally, truster 2 profits if the other truster informs her about the first time that trust is abused, because she can avoid that the trustee takes advantage of her as well.[7]

There exist equilibria comparable to the one for Case 2 if $\tau_{12} \geqslant \mathrm{TEMP}$ and $\tau_{21} < \mathrm{TEMP}$. The main difference is that the trustee now starts randomizing in a period played with

[7] The equilibrium in Case 2 is not unique. The two trusters have to coordinate their randomization probabilities such that the conditions as they are mentioned in the proof are met. However, they still have some freedom how to choose these randomization probabilities. If communicating would be a choice itself, there is another equilibrium in which truster 1 chooses not to communicate information to truster 2. Truster 1 has no incentive to transmit any information to truster 2. Then, both trusters would randomize with a probability TEMP and truster 2 loses the small profits from communication. However, essentially the equilibrium remains the same, and as the last statement in the theorem says, there are no equilibria for which the randomization phase starts considerable later than for the equilibrium described here.

truster 2 and repeats his move in the next period with truster 1 in order to avoid the large information transmission probability of truster 1. Now, the first and last period cause some additional considerations and complicate the formulation of the equilibria without changing the substantive implication that the randomization phase starts when each truster has about the same number of periods left as she would have left in the baseline model.

5. Implications and conclusions

This paper has analyzed effects of adding a third actor to a finitely repeated Trust Game with incomplete information. The main findings of the two-person case remain valid in the three-person games. Trust will be placed in more periods of the game if the truster has less to lose in each period (RISK is smaller). Trust will be placed and honored in the early periods of the game. Thereafter, there is a randomization phase in which eventually trust breaks down and after that no trust will be placed anymore. The probability that trusters place trust in the randomization phase increases with the temptation of the trustee to abuse trust in a given period (TEMP).

This paper provides new results for the three-person cases. First, adding more trustees to the model providing the truster with an exit option does not have an effect on trust. This is a prediction that is compatible with the prediction of many economic models that assume complete information. Here, I have shown that incomplete information and having more types of trustees is not a sufficient condition for exit to be an essential element within the model. Consequently, if effects of exit opportunities are found in experiments, this would imply that the model still lacks some key elements. Second, adding a voice option for the trusters by including two trusters instead of one truster provides more trust only if both trusters inform each other at a high rate about the trustee's behavior. What is considered to be high depends on the temptation for the trustee to abuse trust.

Linking these results to the discussion about learning and control in the beginning of the paper, it can be concluded that an exit option actually does not provide the truster with additional control opportunities in this model. The reason is that the timing of the start of the randomization phase is determined by the payoffs of the truster rather than the payoffs of the trustee. Adding a voice opportunity by introducing a second truster to the model provides the trusters with more learning as well as control possibilities. I hinted in the beginning of this paper at the possibility to determine whether it is more important for a truster to learn from the other truster about the behavior of the trustee or to sanction the trustee by informing the other truster. It turns out that the two aspects of voice need to be combined by the trusters to enable more trust in the trustee. If a truster only exercises control through voice but does not learn from information received from the other truster, control does not have an effect on trust and a truster's trust is based only on her own periods with the trustee. The same is true if the truster only learns from information transmitted by the other truster, but does not control the trustee herself by transmitting information about his behavior.

This last implication of the model provides an opportunity to test the model against models that assume other types of learning than Bayes' updating such as learning by reinforcement (Roth and Erev, 1995; Erev and Roth, 1998). Models based on reinforcement predict that learning has an effect even in the absence of control options for a truster. Although it cannot be excluded that a separate learning effect could also follow from a model with Bayesian updating in which incomplete information is introduced in a different way, such a result would be in favor of a reinforcement model compared to the model developed in this paper. A second type of models that would likely predict effects of exit and more pronounced effects of learning are models in which trusters sometimes experience bad outcomes although the trustee did not intentionally abuse trust. Such situations are expected to be described better by models with imperfect monitoring (see, for example, Radner, 1981; Porter, 1983; Green and Porter, 1984).

A disadvantage of the type of model presented in this paper is that after any abuse of trust, there will never be any doubt about the type of trustee the truster is playing with. As a result, the phase of the game in which there is any learning is limited. Learning is expected to be more important in models in which trustees do not have a fixed type, but there is a small probability that the type of a trustee changes and the trusters cannot observe these changes. Then, trustees are not able to reveal or conceal their type perfectly, so every experience is worthwhile to the trusters (cf. Tadelis, 1999; Mailath and Samuelson, 2001). However, these two models do not consider communication between specified actors. The model in this paper is a first attempt to understand the importance of communication of outcomes among different players in a game. I doubt whether random type changes of the trustee would cause essential changes to the outcomes in my approach, since all (types of) players in the game want to place and honor trust in a considerable part of the game. The randomization phase is the only phase in which different types of trustees act differently, and this is generally only a small number of periods in the game.

The assumption that trusters play in an alternating manner might be changed in, e.g., that one of the trusters is chosen randomly in each period, or that the two trusters play simultaneously with the trustee. Since the equilibria are largely built around the probabilities that information is transmitted between the two rounds, the first change can be analyzed by adapting the probabilities that the "next" truster has information about the last round. The situation that trusters play simultaneously resembles Case 2 with $\tau_{12} = 0$. Consequently, it is still optimal for the trustee to use the same behavior in each round in the two simultaneous games with the trusters. This implies that communication among the trusters cannot improve their outcomes compared to the baseline model. This result as well as the result for Case 2 might change if the trusters receive different payoffs in the game such that in the baseline model, the trustee starts to randomization in a different period with truster 1 than with truster 2 (cf. Bernheim and Whinston, 1990).

Clearly, the discussion about effects of changes in assumptions can be extended much further. However, I think that we lack considerable knowledge about what actually reasonable assumptions are especially related to information availability of actors, information exchange among actors, and how actors actually use this information (belief updating, sanctioning). Especially, in the light of the finding that implications of models such as the finitely repeated Trust Games change dramatically if we move from complete information to incomplete information, is seems necessary to search for (sets of) assumptions under which the findings are more robust to changes that should not have large effects. Although I do not want to start a philosophical discussion at the end of the paper, I would favor a careful design of experiments that do not only allow for testing the implications of theoretical models, but also allow for testing some assumptions especially about how actors use and react to information they obtain in the games they play. Camerer et al. (1993) provide an example for how experiments can be designed in which it is possible to follow more or less the decision making process of the trusters rather than only the final decisions of the actors.

Acknowledgments

Stimulating comments of and discussions with Werner Raub, Chris Snijders, Jeroen Weesie, and an anonymous referee are gratefully acknowledged. Financial support was provided by the Netherlands Organization for Scientific Research (NWO) under grant S 96-168 and PGS 50-370. This paper was largely written during research visits at the Stanford Graduate School of Business and the Department of Sociology at the University of Chicago. I thank both institutes for their support.

Appendix A. Proof of Theorem 4.1

Proof. The theorem is proved by induction on the number of periods each truster plays. First, we reformulate the theorem for $N = 1$.

Beliefs of the trusters

If truster 2 receives information indicating that trust is honored by the trustee *and* $\tau_{12} < \text{TEMP}$, the trustee is definitely friendly ($\pi_1 = 1$). If she receives information about trust honored by the trustee and $\tau_{12} \geqslant \text{TEMP}$, she will update her belief $\pi_1 = \max(\text{RISK}, \pi)$. If she obtains information about abused trust $\pi_1 = 0$.

Strategies of the trusters

If $\tau_{12} < \text{TEMP}$, truster 1 places trust if $\pi \geqslant \text{RISK}$. If $\tau_{12} \geqslant \text{TEMP}$, she places trust if $\pi \geqslant \text{RISK}^2$. Otherwise, truster 1 does not place trust. Truster 2 places trust if $\pi_1 > \text{RISK}$. If $\pi_1 = \text{RISK}$, she places trust with a probability $\min(1, \text{TEMP}/\tau_{12})$. Otherwise, truster 2 does not place trust.

Strategy of a payoff-maximizing trustee

If $\tau_{12} < \text{TEMP}$, a payoff-maximizing trustee always abuses trust. If $\tau_{12} \geqslant \text{TEMP}$, a payoff-maximizing trustee honors trust in period 2 if $\pi \geqslant \text{RISK}$ and honors trust with a probability $(\pi/(1-\pi))(1/\text{RISK} - 1)$ if $\pi < \text{RISK}$. The trustee abuses trust in period 1.

Theorem A.1. *Consider the beliefs and strategies described above.*

- *These beliefs and strategies constitute a sequential equilibrium in the two-period finitely repeated Trust Game with two trusters* ($N = 1$).
- *If* $\pi < \text{RISK}$ *and* $\tau_{12} > \text{TEMP}$, *there is one other sequential equilibrium, namely, never placing trust by both trusters, and always abusing trust by the trustee.*

- *Otherwise, if* $\pi \neq \mathrm{RISK}^n$ *for* $n = 1, 2$, *every sequential equilibrium has the same on-the-equilibrium-path strategies as described before.*

Checking that these beliefs are consistent with Bayes' rationality is straightforward. If the trustee honors trust in period 2 and truster 2 receives this information, then

$$\pi_2 = \frac{\Pr(C_2|\text{friendly})\Pr(\text{friendly})}{\Pr(C_2|\text{friendly})\Pr(\text{friendly}) + \Pr(C_2|\text{payoff-max.})\Pr(\text{payoff-max.})},$$

which results in the given probabilities for all the relevant cases. The only case for which Bayes' rule does not apply is if trust is abused in period 2, while the payoff-maximizing trustee should honor trust with probability 1. The theorem poses that $\pi = 0$ in this out-of-equilibrium situation.

Calculating expected payoffs from the end of the game tree, it is also straightforward to show that the expected payoffs for all three players are optimal given the strategies of the other players. There are no other sequential equilibria because moving backward in the game tree, all the moves are uniquely determined.

Now, I continue with the proof of the main theorem. It can be shown easily that the beliefs of the trusters are consistent with Bayesian updating. In Case 2, the trustee is randomizing only once and uses the outcome of the randomization in any of the two periods in which the truster places trust. Therefore, if truster 1 has placed trust in a pair of moves, she will not obtain any new information if truster 2 informs her after truster 2's period. Moreover, if truster 2 received information in this pair of moves about the behavior of the trustee in the previous period, she will not update again her beliefs after her own period.

Using the induction argument, it has to be proved for both cases that the first move of each truster and the first two moves by the trustee are optimal. If both trusters have no incentive to withhold trust in period $2N-2$ and $2N-3$, the trustee will honor trust in the first two periods: he obtains two times R_2 and, thereafter, he is in a similar position as before. If he abuses trust immediately, he will shift to P_2 payoffs thereafter, losing the opportunity to receive some R_2 payoffs. Consequently, the trusters will place trust in their first periods of play if they are going to place trust in periods $2N-2$ and $2N-3$. Thus, the key cases to be checked are those where the trusters probably do not place trust in their first or second period of play.

Case 1. Since the trustee's strategy is exactly the same as in Theorem 2.1 (he acts as if he is playing $2N$ periods with one truster), the trusters play an optimal response, because they behave as if only one truster is involved with the only exception that they use other randomization probabilities. Now, we check optimality of the trustee's behavior.

Assume $\mathrm{RISK}^{2N-2} < \pi < \mathrm{RISK}^{2N-3}$. From the induction assumption, it is known that the trustee is indifferent in period $2N-2$ and will be randomizing in this period. For the calculation of the expected payoffs, it can be assumed that he abuses trust in this period and in period $2N-3$.[8] Now, consider the first two moves of the trustee.

$$\begin{aligned}
u_2(C_2C_2) &= R_2 + R_2 + T_2 + \tau_{12}P_2 + (1-\tau_{12})\big(p_{2N-3}T_2 + (1-p_{2N-3})P_2\big) + (2N-4)P_2 \\
&= 2R_2 + T_2 + (2N-3)P_2, \\
u_2(C_2D_2) &= R_2 + T_2 + \tau_{21}P_2 + (1-\tau_{21})T_2 + (2N-3)P_2 < 2R_2 + T_2 + (2N-3)P_2, \\
u_2(D_2C_2) &= T_2 + \tau_{12}P_2 + (1-\tau_{12})R_2 + (2N-2)P_2 < u_2(D_2D_2), \\
u_2(D_2D_2) &= T_2 + \tau_{12}P_2 + (1-\tau_{12})T_2 + (2N-2)P_2 < R_2 + T_2 + (2N-2)P_2.
\end{aligned}$$

Consequently, two times honoring trust is equilibrium play for the trustee, which is the behavior specified in the theorem. In a similar vein, it can be shown that the trustee is indifferent between C_2C_2 and C_2D_2 and prefers these two combinations over D_2C_2 and D_2D_2 if $\mathrm{RISK}^{2N-1} < \pi < \mathrm{RISK}^{2N-2}$. Finally, if $\pi < \mathrm{RISK}^{2N-1}$, the

[8] Throughout this proof I use this argument to calculate the payoffs in the periods that are covered by the induction assumption, which mostly implies that payoffs are calculated for the trustee abusing trust and truster not placing trust, because then the payoffs for the remaining periods can be determined straightforwardly.

trustee turns out to be indifferent between any combination of moves to start with in the first two rounds. The existence of a second equilibrium in this situation follows from the same considerations as given for $N = 1$. That there are no other equilibria follows from the uniqueness of the corresponding equilibrium in the baseline model.

Case 2. If the trustee honors trust with certainty, the trusters place trust in period $2N$ and $2N - 1$. Therefore, I only need to consider the case were the trustee randomizes in period $2N$, i.e., if $\pi_{2N} < \mathrm{RISK}^{N-1}$ in period $2N$ with truster 1. Then, it holds for truster 1 that

$$\begin{aligned} u_1(C_1) &= \pi R_1 + (1-\pi)\bigl(q_{2N} R_1 + (1-q_{2N})S_1\bigr) + (N-1)P_1, \\ u_1(D_1) &= N P_1 \end{aligned}$$

implying that $u_1(C_1) > u_1(D_1) \Leftrightarrow \pi_{2N} > \mathrm{RISK}^N$. Thus, truster 1 is acting optimal. The expected payoffs for truster 2 if truster 1 did not place trust or if truster 2 did not receive information about the trustee's behavior in period $2N$ are

$$\begin{aligned} u_1(C_1) &= \pi R_1 + (1-\pi)\bigl(q_{2N} R_1 + (1-q_{2N})S_1\bigr) + u_1(\text{after } C_1), \\ u_1(D_1) &= P_1 + u_1(\text{after } D_1), \end{aligned}$$

where $u_1(\text{after } C_1) = u_1(\text{after } D_1)$, because the behavior of truster 1 does not depend on whether truster 2 informs her about the trustworthiness of the trustee. If truster 1 did not place trust once, she will never place trust again although the trustee might still be trustworthy toward truster 2. If truster 1 did place trust, she knows what the trustee will play in the following period with truster 2, so nothing will change for her whether or not truster 2 informs her before her next period. Therefore, the probability that truster 2 obtains information about the trustee in a foregoing period does not depend on her own behavior. Consequently, $u_1(C_1) > u_1(D_1) \Leftrightarrow \pi_{2N} > \mathrm{RISK}^N$, and truster 2 is indifferent if equality holds.

Now consider the four possible strategies for the trustee for the first two periods. Again I only need to consider the situation in which the trusters probably do not place trust in their second period, i.e., if $\pi_{2N} < \mathrm{RISK}^{N-1}$.

Before the calculation, note that for the randomization probabilities of the trusters holds

$$p_{2N-2} + p_{2N-3} - \tau_{12} p_{2N-2} p_{2N-3} = 2\mathrm{TEMP} - \tau_{12}, \tag{A.1}$$

$$0 < p_{2N-3} = \frac{\mathrm{TEMP}(1-\tau_{12}) - \tau_{12}\mathrm{TEMP} + \tau_{12}^2}{(1-\tau_{12}) - \tau_{12}\mathrm{TEMP} + \tau_{12}^2} < \frac{\mathrm{TEMP}}{1-\tau_{12}}, \tag{A.2}$$

$$0 < p_{2N-2} = \frac{\mathrm{TEMP} - \tau_{12}}{1-\tau_{12}} < \frac{\mathrm{TEMP} - \tau_{12}}{1-\tau_{21}}. \tag{A.3}$$

Condition (A.1) follows from straightforward manipulation. For Condition (A.2), one has to realize that the enumerator equals $\mathrm{TEMP} - \mathrm{TEMP}^2 + (\mathrm{TEMP} - \tau_{12})^2 > 0$ and that $-\tau_{12}\mathrm{TEMP} + \tau_{12}^2 < 0$. Condition (A.3) follows from $\tau_{12} < \tau_{21}$.

Now, it can be shown that the trustee is indifferent between playing D_2D_2 and C_2C_2, while C_2D_2 and D_2C_2 provide him with a lower payoff.

$$\begin{aligned} u_2(C_2C_2) &= R_2 + R_2 + p_{2N-2}T_2 + (1-p_{2N-2})P_2 + p_{2N-2}\tau_{12}P_2 \\ &\quad + (1-p_{2N-2}\tau_{12})\bigl(p_{2N-3}T_2 + (1-p_{2N-3})P_2\bigr) + (2N-4)P_2 \\ &= 2R_2 + \bigl(p_{2N-2} + p_{2N-3} - \tau_{12}p_{2N-2}p_{2N-3}\bigr)(T_2 - P_2) + (2N-2)P_2 \\ &= 2T_2 - \tau_{12}(T_2 - P_2) + (2N-2)P_2, \\ u_2(C_2D_2) &= R_2 + T_2 + \tau_{21}P_2 + (1-\tau_{21})\bigl(p_{2N-2}T_2 + (1-p_{2N-2})P_2\bigr) + (2N-3)P_2 \\ &= R_2 + T_2 + p_{2N-2}(1-\tau_{21})(T_2 - P_2) + (2N-2)P_2 \\ &< R_2 + T_2 + (\mathrm{TEMP} - \tau_{12})(T_2 - P_2) + (2N-2)P_2 \end{aligned}$$

$$
\begin{aligned}
&= 2T_2 - \tau_{12}(T_2 - P_2) + (2N-2)P_2,\\
u_2(D_2C_2) &= T_2 + \tau_{12}P_2 + (1-\tau_{12})R_2 + P_2 + (1-\tau_{12})^2\bigl(p_{2N-3}T_2 + (1-p_{2N-3})P_2\bigr)\\
&\quad + \bigl(1-(1-\tau_{12})^2\bigr)P_2 + (2N-4)P_2\\
&= T_2 + R_2 - \tau_{12}(R_2 - P_2) + (1-\tau_{12})^2 p_{2n-3}(T_2 - P_2) + (2N-2)P_2\\
&< T_2 + R_2 - \tau_{12}(T_2 - P_2) + \tau_{12}(T_2 - R_2) + (1-\tau_{12})(T_2 - R_2) + (2N-2)P_2\\
&= 2T_2 - \tau_{12}(T_2 - P_2) + (2N-2)P_2,\\
u_2(D_2D_2) &= T_2 + \tau_{12}P_2 + (1-\tau_{12})T_2 + (2N-2)P_2 = 2T_2 - \tau_{12}(T_2 - P_2) + (2N-2)P_2.
\end{aligned}
$$

Consequently, $u_2(C_2C_2) = u_2(D_2D_2)$ and it is indeed optimal for the trustee to randomize between these two pairs of moves. It is easy to check that if truster 1 does not place trust in her first period, the trustee has still no incentive do deviate from his randomization strategy.

There are no equilibria in which the trustee starts randomizing later. It can be checked easily that for $N = 2$ and $\pi <$ RISK, the trustee has to start randomizing or abusing trust at least in his first encounter with truster 2, which is period 3. Using again an induction argument it follows that the trustee cannot continue to place trust with certainty beyond an encounter with truster 2 for which $\pi < \text{RISK}^n$ in period $(n-1)/2$. □

References

Berg, J., Dickhaut, J., McCabe, K., 1995. Trust, reciprocity, and social history. Games Econ. Behav. 10, 122–142.

Bernheim, B.D., Whinston, M.D., 1990. Multimarket contact and collusive behavior. RAND J. Econ. 21, 1–26.

Bower, A.G., Garber, S., Watson, J.C., 1997. Learning about a population of agents and the evolution of trust and cooperation. Int. J. Ind. Organ. 15, 165–190.

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

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

Camerer, C.F., Johnson, E.J., Rymon, T., Sen, S., 1993. Cognition and framing in sequential bargaining for gains and losses. In: Binmore, K., Kirman, A., Tani, P. (Eds.), Frontiers of Game Theory. MIT Press, Cambridge, MA, pp. 27–47.

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

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. Amer. Econ. Rev. 88, 848–881.

Green, E.J., Porter, R.H., 1984. Noncooperative collusion under imperfect price information. Econometrica 52, 87–100.

Güth, W., Kliemt, H., 1994. Competition or co-operation: on the evolutionary economics of trust, exploitation and moral attitudes. Metroecon. 45, 155–178.

Güth, W., Ockenfels, P., Wendel, M., 1997. Cooperation based on trust: an experimental investigation. J. Econ. Perspect. 18, 18–43.

Hirschman, A.O., 1970. Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States. Harvard Univ. Press, Cambridge, MA.

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

Kreps, D.M., Milgrom, P., Roberts, J., Wilson, R., 1982. Rational cooperation in the finitely repeated prisoners' dilemma. J. Econ. Theory 27, 245–252.

Kreps, D.M., Wilson, R., 1982a. Reputation and imperfect information. J. Econ. Theory 27, 253–279.

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

Lahno, B., 1995. Trust, reputation, and exit in exchange relationships. J. Conflict Resolution 39, 495–510.

Macy, M.W., Skvoretz, J., 1998. The evolution of trust and cooperation between strangers: a computational model. Amer. Soc. Rev. 63, 638–660.

Mailath, G.J., Samuelson, L., 2001. Who wants a good reputation? Rev. Econ. Stud. 68, 415–441.

McKelvey, R.D., Palfrey, T.R., 1992. An experimental study of the centipede game. Econometrica 60, 803–836.

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

Porter, R.H., 1983. Optimal cartel trigger-price strategies. J. Econ. Theory 29, 313–338.

Radner, R., 1981. Monitoring cooperative agreements in a principal-agent relationship. Econometrica 53, 1173–1198.

Raub, W., Weesie, J., 1990. Reputation and efficiency in social interactions: an example of network effects. Amer. J. Soc. 96, 626–654.

Raub, W., Weesie, J., 2000. The management of matches: a research program on solidarity in durable social relations. Netherlands' J. Soc. Sci. 36, 71–88.

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

Schüßler, R.A., 1989. Exit threats and cooperation under anonymity. J. Conflict Resolution 33, 728–749.

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

Snijders, C., Keren, G., 1999. Determinants of trust. In: Budescu, D.V., Erev, I., Zwick, R. (Eds.), Games and Human Behavior. Lawrence Erlbaum, Mahwah, NJ, pp. 355–385.

Tadelis, S., 1999. What's in a name? Reputation as a tradeable asset. Amer. Econ. Rev. 89, 548–563.

Vanberg, V., Congleton, R.D., 1992. Rationality, morality, and exit. Amer. Polit. Sci. Rev. 86, 418–431.

Weesie, J., 1996. Disciplining via exit and voice, ISCORE Paper No. 88, Utrecht University.