



Research Paper

On the relational aspects of trust and trustworthiness: Results from a laboratory experiment

Federica Alberti^a, Anna Conte^b, Daniela T. Di Cagno^c, Emanuela Sciubba^{d,*}

^a Faculty of Business and Law, University of Portsmouth, UK

^b Department of Statistical Sciences, Sapienza University of Rome, Italy

^c Department of Economics and Finance, LUISS Guido Carli, Rome, Italy

^d Birkbeck Business School, London, UK



ARTICLE INFO

JEL classification:

C72
C91
C92
D82
D85

Keywords:

Social network
Trust
Trustworthiness
Lab experiment

ABSTRACT

Do we trust better-connected people more than others and are those who are better connected more trustworthy? Interaction in social networks affects trust as it helps reduce informational asymmetries and identify those whom to trust. It also provides channels for reciprocity where trustworthiness emerges as a relational rather than individual characteristic. We run a laboratory experiment in which trustors decide on how to allocate their endowment to five trustees both before and after a network formation phase. Our results show that trustors are influenced by the behaviour of trustees in the network formation phase: when the allocation is chosen after networks have been formed, trust is directed towards the most frequent connections of the trustor. However, when trustors' offers are anonymous, such increased trust is not reciprocated by trustees. This suggests that social connections do not signal a greater individual propensity to reciprocate trust, but rather provide channels for reciprocity which can only be activated when both trust and trustworthiness are not anonymous.

1. Introduction

Trust and trustworthiness are often seen as two sides of the same coin, while they have different determinants and motivations. Both are rooted in social interaction and past experience, but not necessarily through the same channels. Interactions in social networks affect trust as they help reduce informational asymmetries and identify those whom to trust. They also provide a structure for repeated interaction and reciprocity where trustworthiness emerges as a relational rather than individual characteristic.

Since the seminal contribution by Berg et al. (1995), both trust and trustworthiness are seen as the result of an individual and a social, or relational, component.¹ Trust, in particular, builds on an innate orientation towards others, rooted in social preferences such as altruism or generosity and a component of calculated risk-taking, which Torbiörn (2002) defines as “the willingness to be

* Corresponding author.

E-mail addresses: federica.alberti@port.ac.uk (F. Alberti), anna.conte@uniroma1.it (A. Conte), ddicagno@luiss.it (D.T. Di Cagno), e.sciubba@bbk.ac.uk (E. Sciubba).

¹ A large body of experimental literature explores trust and trustworthiness (see Camerer, 2003, for a review) and unravels the psychological, social and ethical motivations involved in trusting relationships (see Ostrom and Walker, 2003).

<https://doi.org/10.1016/j.jebo.2023.12.031>

Received 17 November 2022; Received in revised form 5 December 2023; Accepted 27 December 2023

Available online 26 January 2024

0167-2681/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

vulnerable” that reflects the expectation of trustworthiness or reciprocation by others. While altruism is a personal trait, “expectation of trustworthiness” (Barr, 2003; Ashraf et al., 2006; Salgado et al., 2021) is built through experience in a social network (see Granovetter, 1985). Tullberg (2008) distinguishes between personal traits versus experience as main drives of trust and claims that trust increases with the degree of perceived trustworthiness in society.

Trustworthiness has an individual and a social or relational component too. Tullberg (2008) finds that individuals are more inclined to be trustworthy towards those who are seen as trusting others, independently from their individual attitude to donate. Bicchieri et al. (2011) argue that trust should be seen more as an individual characteristic while trustworthiness more as a social norm and Putnam (2000) considers trustworthiness as more important than trust for social capital or for what Jacobs (1961) calls ‘a neighbourhood network’.

In this paper, we propose an experimental design that explores the individual and the relational components of trust and trustworthiness. We do so by having subjects play a trust game both before and after a network formation game which mimics social interaction in a meaningful and payoff-relevant way. In the network formation game subjects need to rely on each other and coordinate in order to invest in costly mutual links which generate higher payoffs when they provide access to a larger network of social connections.

A number of studies have combined field and lab data to study the impact of social networks on trust and trustworthiness. Glaeser et al. (2000) measure trust and trustworthiness among a sample of Harvard undergraduates. They find that social connections matter but only to the extent that better-connected subjects attract higher returns in trust games. Riyanto and Jonathan (2018) match data from subjects’ positions in online social networks to their behaviour in a laboratory trust game. They find that, while subjects tend to trust recipients who are well-connected and have a greater centrality in the social network, central individuals are not necessarily more trustworthy, with both trust and trustworthiness being affected by closeness and proximity in the social network.

Di Cagno and Sciubba (2010) conduct an experiment where networks are generated endogenously in the laboratory and compare the behaviour in a trust game when this proceeds or follows the network formation phase to when it is played on its own. Akin to the lab-in-field experiments (Glaeser et al., 2000 and Riyanto and Jonathan, 2018, among others), they find that subjects trust more those who are better connected or who are closer to them in the social network. However, unlike the lab-in-field experiments, they find that trustworthiness is significantly higher when the trust game is played after a network formation phase.

The positive impact of endogenous network formation on trustworthiness can be explained by two separate drives, which could be operating concurrently. By interacting in a social network and observing how their potential trustees behave in the network formation phase, trustors manage to direct their trust towards more trustworthy recipients. Hence, when trustworthiness is an individual unobservable characteristic, interaction in social networks helps alleviate informational asymmetries and allows trustors to direct their trust more efficiently. At the same time, such enhanced trustworthiness could be explained by the fact that it is a relational, rather than individual characteristic: trustees are more trustworthy towards those they have interacted with positively during the network formation phase, essentially reciprocating with their responses both the generosity of the offers received in the trust game and the behaviour observed and experienced in the network formation phase.

In this paper, we modify the experimental design in Di Cagno and Sciubba (2010) with the aim of investigating whether, in the absence of knowledge of the identity of the trustor (relational component), trustworthiness still holds.

We design our experiment around an involuntary trust game (McCabe et al., 2003) in which the trustor decides how to allocate his/her endowment across a number of trustees, with no option to retain the remainder.² The trustor’s behaviour is therefore no longer affected by his/her altruism or generosity, but only by his/her directed trust or expected reciprocity from each of the trustees. Throughout, we will address this game as the “trust-allocation” game.

We let participants play a trust-allocation game before and after a network formation game. We find that, when trust is allocated before networks are formed, endowments are allocated somewhat uniformly across trustees. When trust is instead allocated after networks are formed, trustors direct their trust towards a smaller number of trustees. In particular, trust increases with the number of times the trustor and the trustee have been directly linked, and decreases with the number of times they have been indirectly linked, or if the trustee has been isolated. Being well-connected via indirect links attracts lower trust, and having borne the cost of many direct links does not seem to matter greatly.

For our measurement of trustworthiness, we elicit trustees’ responses by using the strategy vector method (Selten, 1965). Unlike Di Cagno and Sciubba (2010), we only condition the trustees’ responses on the amount received and the characteristic of the trustee, but not also on the identity of the trustor. This implies that, when the trust-allocation game is played after the network formation game, trustees are not able to reciprocate trustors’ behaviour in the network formation phase, but only trustors’ offers in the trust-allocation game. Hence responses reflect only their individual, and not also their relational, or social, trustworthiness.

We compare trustees’ responses before and after the network formation game. We find that, when trustors’ identities are not disclosed to trustees, the trust allocation made after social interaction does not receive a higher response than their initial allocation of trust. In fact, returns to trust are actually lower after the network formation game, as well as the trustor’s earnings. We attribute the decrease in trustworthiness to trustees’ expectations that trustors will concentrate trust on a smaller number of trustees after the network formation game. To remove the effect of trustees’ expectations, we compare trustor’s earnings before the network formation game to what trustor’s earnings would have been had they allocated trust as they did after the network formation game but with the trustees responding as they did before the network formation game. We find no differences in trustors’ earnings here. Finally, we estimate econometric models of trustworthiness against characteristics of the trustee, including the number of times the trustee

² Involuntary trust games arise in many real-life situations in which we are obliged to find someone to trust.

had been isolated in the network and the number of both trustee’s direct and indirect links. We find no effects of these variables on trustworthiness. This implies that the social network is not informative of individual trustworthiness.

This indicates that the higher level of trustworthiness among network friends or partners, as observed in the field by Glaeser et al. (2000)³ and in the lab by Di Cagno and Sciubba (2010), is not due to the ‘information value’ of social interaction, which would allow trustors to direct their trust towards more trustworthy recipients, but rather to the reciprocal behaviour that is induced by the relationship between trustor and trustee established during the network formation phase.

The paper develops as follows: in Section 2 we introduce a theoretical framework; in Section 3 we describe the experimental design; Section 4 presents and discusses our main results and Section 5 concludes.

2. A model of trust and trustworthiness among socially connected individuals

In the standard version of the trust game (Berg et al., 1995), two players, a trustor and a trustee, are randomly and anonymously matched to play the following game: the trustor (he/his) is endowed with a monetary amount $m > 0$ and has to decide which share of his endowment, if at all, to offer to a trustee (she/her). If the trustor offers amount o to the trustee, with $0 \leq o \leq m$, the trustee receives a multiple of o , usually $3o$, and has to decide which share of the amount received, if at all, she wishes to return to the trustor. If the trustee return r , with $0 \leq r \leq 3o$, payoffs for the trustor and the trustee are respectively:

$$m - o + r$$

$$3o - r$$

The unique rational solution for selfish trustor and trustee is $(o^*, r^*) = (0, 0)$, with payoffs m and 0 for the trustor and the trustee respectively: the trustor anticipates that the trustee will return nothing and keeps the entire endowment for himself. This outcome is not Pareto efficient as both the trustor and the trustee can be better off with a positive offer from the trustor, as long as the trustee is trustworthy and returns to the trustor an amount at least equal to the offer received. We can define the trustee’s trustworthiness as $t = \frac{r}{o}$, with $t \geq 1$ and $o > 0$ resulting in a Pareto improvement.

In the real world, we rarely play trust games with complete strangers. More often we choose whom to trust among those whom we are socially connected to and we have some expectations on the trustworthiness of our trustees which are rooted in previous interactions we may have had with them.

In this paper, we adapt the trust game to a multi-player setting in which trustors choose how to allocate their trust among several trustees. Also, we allow for the possibility that players are known to each other for their previous social interaction, which may provide information to trustor on trustees’ trustworthiness and allow for greater returns to trust.

We consider a finite number of trustors indexed by $i = 1, 2, \dots, I$ and of trustees indexed by $j = 1, 2, \dots, J$. Each trustor i has to decide how much to offer to each of the trustees and we denote by o_{ij} the offer that trustor i makes to trustee j , with $o_{ij} \geq 0$ and $\sum_{j=1}^J o_{ij} \leq m$. Similarly to the 2-players version of the game, trustee j receives $3o_{ij}$ from trustor i and has to decide on the amount to return to i which we denote r_{ji} , with $0 \leq r_{ji} \leq 3o_{ij}$. We interpret the ratio $t_{ji} = \frac{r_{ji}}{o_{ij}}$ as the trustworthiness of trustee j towards trustor i and we note that when $t_{ji} \geq 1$ positive offers result in a Pareto improvement with respect to the selfish solution.

In a multi-player setting, a natural question to ask is whether trust and trustworthiness are individual traits of the trustor and trustee and/or relational characteristics which depend on both subjects involved.

Let us start from trust. A trustor may make a larger offer because he is altruistic. The trustor’s social preferences are an individual characteristic which will result in larger offers irrespectively of the identity of the trustee. At the same time a trustor may make a larger offer to a particular trustee because he expects the trustee to be trustworthy, or because he is reciprocating with greater trust positive behaviour that a particular trustee showed towards him or towards others. This may make a trustor more trusting towards some trustees than towards others.

In this paper, we consider an involuntary trust game *à la* McCabe et al. (2003) in which the trustor has to allocate his entire endowment and cannot keep any of it for himself, so that $\sum_{j=1}^J o_{ij} = m$. This allows to abstract from the general level of altruism of the trustor (how selfless he is) and to focus entirely on his trust allocation strategy. A trustor will allocate trust to those trustees who he expects will return more money; also, a trustor who is a reciprocator may wish to reward with greater trust those who have expressed positive behaviour towards him or towards others. For the purpose of this paper, we will call *perceived trustworthiness* whatever makes a trustee more deserving of trust from the point of view of the trustor.

Given trustor i ’s perceived trustworthiness of trustees $j = 1, 2, \dots, J$, $\hat{t}_i = (\hat{t}_{1i}, \hat{t}_{2i}, \dots, \hat{t}_{Ji})$, the optimal trust allocation would be to entrust his entire endowment⁴ to trustee j^* such that

$$j^* = \arg \max_j \hat{t}_{ji}$$

where \hat{t}_{ji} is trustor i ’s perceived trustworthiness of trustee j .

³ See also Leider et al. (2009) and Binzel and Fehr (2013).

⁴ In case of a tie, if there are many trustees with the same highest level of trustworthiness, the trustor will be indifferent and may allocate to only one of them, or to any combination of them.

For a selfish trustor i , perceived trustworthiness of trustee j coincides with j 's trustworthiness, i.e. the share of the offer received by i that j returns:

$$\hat{t}_{ji} = t_{ji} = \frac{r_{ji}}{o_{ij}}$$

If the trustor is a reciprocator, i 's assessment of j 's trustworthiness will also be affected by a relational component $k_{ij} \geq 0$ which may amplify (when $k_{ij} > 1$) or reduce (when $k_{ij} < 1$) i 's perceived trustworthiness of j :

$$\hat{t}_{ji} = k_{ij}t_{ji} = k_{ij} \frac{r_{ji}}{o_{ij}}$$

If trustworthiness is not observable, a rational trustor should allocate his endowment among trustees whom he perceives as more likely to be trustworthy.

We usually choose whom to trust relying on our assessment of trustees' trustworthiness acquired through what we know about them and in particular through what we have observed or experienced in any interaction we have previously had with them.

Here we term all such interactions *social interaction* and we assume that, through social interaction, each trustor receives a signal which is informative about the trustworthiness of each of the trustees, to include the likelihood which they will return a high share of what they have received, and any other trait that make them appear more deserving, for example because they exhibited kind behaviour towards the trustor or towards others. Before receiving the signal, the trustor perceives all trustees to be trustworthy with the same probability $p > 0$. Therefore, in the absence of informative signals, a rational trustor will allocate trust uniformly across all the trustees.

When a signal on trustees' trustworthiness is received, the trustor will update his beliefs and will deem some trustees to be more trustworthy than others. In particular, engaging in social interaction will provide trustor i with a binary signal $s_j \in \{0, 1\}$ on j 's trustworthiness. If signal $s_j = 1$ is observed, then $\Pr(t_{ji} \geq 1 \mid s_j = 1) > p$; while if signal $s_j = 0$ is observed, then $\Pr(t_{ji} \geq 1 \mid s_j = 0) < p$, which implies that the allocation of trust after the social interaction will be concentrated on those trustees for which a positive signal has been observed.

The latter provides a testable implication of the relational component of trust which will be directed toward those trustees who are perceived to be more trustworthy.

Next, we move to trustworthiness. Similarly to trust, trustworthiness can be seen as the result of two components: an innate component that is hardwired to the individual and a relational component whereby the individual will display different levels of trustworthiness to different trustors.

In order to distinguish between the individual and the relational components of trustworthiness, we assume that:

$$t_{ji} = k_{ji}t_j$$

where t_j is the innate disposition of subject j to reciprocate trust and $k_{ji} \geq 0$ is a factor which depends on the relationship between trustee j and trustor i which may amplify (when $k_{ji} > 1$) or reduce (when $k_{ji} < 1$) innate reciprocity.

To simplify, we assume that t_{ji} can take only one of two values: trustees either return the amount received ($t_{ji} = 1$) or none ($t_{ji} = 0$). Let us call trustworthy the trustee that returns the amount received ($t_{ji} = 1$) and untrustworthy the one that does not ($t_{ji} = 0$). Furthermore, we assume that $t_j \in \{0, 1\}$ and that $k_{ji} \in \{0, 1\}$, so that a trustee is trustworthy when she is both innately so ($t_j = 1$) and she has a positive relation with trustor i ($k_{ji} = 1$).

From the point of view of trustor i , perceived trustworthiness of trustee j depends on j 's trustworthiness towards i , t_{ji} , but also on the relational component $k_{ij} \in \{0, 1\}$, so that trustor i perceives trustee j trustworthy when she is both trustworthy towards him ($t_{ji} = 1$) and he has a positive relationship with her ($k_{ij} = 1$):

$$\hat{t}_{ji} = k_{ij}t_{ji} = k_{ij}k_{ji}t_j$$

Hence trust allocation is driven by perceived trustworthiness which in turn is potentially the result of three separate components: the first one (k_{ij}) has to do with how deserving (of trust) trustee j appears to trustor i ; the second one (k_{ji}) has to do with how deserving (of trustworthiness) trustor i appears to trustee j ; and finally the third component is the innate trustworthiness of trustee j (t_j).

We compare trust-allocation strategies when trustors allocate trust across strangers, or rather interact with trustees in a network formation game before making their trust-allocation decisions.

By observing trustors' allocation strategies, trustworthiness and returns to trust before and after the network formation game, we are able to test the following three hypotheses:

H1: *Trustors condition their trust-allocation strategy on information acquired through social networks: we expect trustors to allocate more trust to those for which positive signals have been observed during the social interaction phase and are therefore perceived to be more trustworthy.*

H2: *Trustees who do not observe the identity of the trustor respond to trust on the basis of individual trustworthiness alone. Hence, we do not expect trustworthiness to differ before or after social interaction.*

H3: *Social networks are informative about trustworthiness. We expect the trustor's earnings to be greater after social interaction.*

3. Experimental design

We designed a sequential experiment composed of three parts, which alternate a trust-allocation game and a network formation game.

In part 1, all subjects played both roles, i.e. trustor and trustee, in a trust-allocation game where responses were elicited through the strategy vector method (Selten, 1965). In part 2, participants played at least 30 rounds of a network formation game. Part 3 is the same as part 1 and subjects played both roles in a trust-allocation game, again with responses elicited through the strategy vector method.

We describe the designs for the network formation game and the trust-allocation game in subsections 3.1. and 3.2, while below we focus on the overall design and the experimental procedures.

The experiment was run at the Experimental Economics Laboratory of the Max Planck Institute of Economics in Jena (Germany). Participants were mainly students from various backgrounds at the Friedrich Schiller University. Participants were recruited using ORSEE (Greiner, 2015), and each participant attended only one session.

Upon arrival, participants were randomly allocated to cubicles, where they sat in front of the computer. Instructions were read aloud, and time was given to answer control questions. The experiment was computerized using z-Tree (Fischbacher, 2007). Subjects were not allowed to communicate with each other during the experiment except for interacting via computers.

Participants were randomly matched by the computer in groups of six subjects at the beginning of each session, and group composition was kept fixed until the end of the session.

At the beginning of the session, participants were informed that the experiment comprised three parts and that they would receive instructions for each part at the outset of the relevant part.⁵ Subjects were encouraged to ask for clarification after reading the instructions and at any time during the experiment.

We ran 14 sessions with 30 participants per session, involving a total of 420 participants and 70 groups overall.

We adopted a within-subjects design with a single treatment. Total earnings from the experiment expressed in Experimental Currency Units (ECUs) were defined by the earnings in one of the two trust-allocation games in part 1 and part 3, plus the earnings from one of the rounds of the network formation game in part 2.

At the end of each session and for each of the trust-allocation games, the computer randomly drew one of the six participants as ‘trustor’ and his trust-allocation decisions were applied to the other five players in the group, which were consequently selected as trustees, with their return decisions retrieved and applied.

To determine which of the two trust-allocation games was relevant for earnings, in each session one of the participants was randomly selected and asked to draw a ball from an urn containing two balls, one labelled with the number 1 and another labelled with the number 3; the selected ball determined which part of the experiment (either part 1 or part 3) was relevant for participants’ earnings. The earnings from part 2 (network formation game) were determined by the earnings of one, randomly selected, round played by the participants. This payment structure mitigated learning effects across the two trust-allocation games (part 1 and part 3) and income effects across the three experimental parts and across rounds of the network formation game (part 2).

For the trust-allocation games, the conversion rate was 1 ECU = €0.80; for the network formation game, the conversion rate was 1 ECU = €0.02. The average earnings were about €19. In addition to the monies earned in the experiment, all participants received a show-up fee of €2.50. Payments were made individually and anonymously at the end of each session.

This experimental design allows us to measure the impact of a social network on trust and trustworthiness. By comparing trust-allocation decisions in part 1 and part 3, we are able to identify the effect of social connections on trust and in particular how interaction in a social network may signal some subjects as more deserving of trust. We then go on to assess if such directed trust is well placed on the most trustworthy individuals. In order to elicit individual trustworthiness, we maintain trustors’ anonymity and enable trustees to condition their responses solely on the amounts received, but not on the identity of the trustor. This approach controls for any inclination trustees may have to reciprocate trustors for their participation in the network formation phase, which reflects what we refer to as social or relational, rather than individual, trustworthiness.

Di Cagno and Scuibba (2010) run a between-subjects laboratory experiment in which subjects play a trust game by selecting their trustees after having interacted with them in a network formation game. They find that trustors are more likely to select as trustee someone they had positive social interaction with (because they established repeated links with them, for example). They also find that the level of trustworthiness and returns to trust are increased when trustors choose their trustees after the network formation game, however, the overall level of trust (size of the offers) is lower. We argue that this potentially confounding result is because their design does not allow to separate between the individual and the relational components of trust and trustworthiness. In particular, for the trustor, their design only allows to test jointly for where the trust is directed (the choice of whom to trust) and the level of trust (the choice of how much to offer). Also, given that trustees observe both the offer received and trustors’ identities, one cannot be sure whether the higher returns to trust are due to the fact that the network formation phase is successful at signalling higher levels of innate trustworthiness (and therefore a more generous response to the offer received), or rather due to the relational component of trustworthiness (reciprocating behaviour after the network formation game).

In this paper, for part 2 of the experiment, we employ the same network formation game as in Di Cagno and Scuibba (2010). However, our designs differ substantially in many respects.

⁵ Instructions are provided in Appendix A. The original instructions in German are available upon request.

The most obvious difference is that in parts 1 and 3 we use trust-allocation games rather than traditional trust games. This enables us to disentangle the separate issues of whom to trust from how much to trust. Additionally, we adopt a within-subjects design in which trust-allocation game responses are elicited via the strategy-vector method. With this, we aim to compare returns to trust in part 1 to what the returns to trust would have been if trustors had allocated trust as in part 3 (with trustees' responses as in part 1), and therefore isolate the effect of the social interaction *ceteris paribus* as far as income or learning effects are concerned. Finally, in our design, trustees' responses are conditional on the amount received but not also on the trustors' identities. Therefore, they reveal how much each trustee is willing to reciprocate a trustor's offer (individual trustworthiness), but not also how much trustees are willing to reciprocate trustors for the interaction they have had with them in the network formation game (relational trustworthiness).

The next two sections describe the network formation and the trust-allocation games in more detail.

3.1. The network formation game

We follow Di Cagno and Sciubba (2010) for the experimental implementation of the network formation game.⁶ Players can propose links to one another. There are two types of links: direct and indirect. A direct link implies that two players are mutually proposing a link. An indirect link is a link to other players accessed via an existing link (either direct or indirect). If subject A is directly linked to subject B and subject B is directly linked to subject C, subjects A and C are indirectly linked. A player can be isolated by having no links – either direct or indirect – with anyone. However, a player can earn only by linking her/himself to others. Direct and indirect links are equally beneficial. Conversely, direct links are costly, unlike indirect links that are free to establish. Therefore, a profit-maximising player will seek to establish as many indirect links as possible.

The network formation game involves six players and consists of a minimum of 30 rounds with a random stopping rule after the thirtieth round.⁷ At the beginning of each round, all players are isolated (i.e. there are no formed links among players within a group) and they are asked whether they want to propose any link to the other participants and, in case, to whom. Each player can propose up to 5 links. All players submit link proposals simultaneously. The computer collects the proposals from all participants, and displays the activated links on the screen by means of a line which connects the players who are linked. Each player is identified by an icon, that is @, #, ¶, ±, ~, or ÷.⁸ The symbol representing a player is inscribed in a blue circle.

At the beginning of each round, the computer assigns to each player an initial endowment of 450 ECUs which is equal across participants. Each link gives 100 ECUs, while the cost of a direct link is 90 ECUs. The profit rule is:

$$\text{Profit} = 450 + 100 \times \text{number of participants reached (directly and indirectly)} \\ - 90 \times \text{number of direct links}$$

Other features of the experimental version of the game are:

- Profits are displayed on the computer screen;
- The screen provides further information:
 - whether the players have received link proposals (and by which players) that did not result in a direct link because a link to those players was not proposed; and
 - whether the players have made link proposals that have not been accepted.⁹

3.2. The trust-allocation game

The trust-allocation game involves six players, identified by the same icons used in the network formation game. Two identical trust-allocation games are played one-shot, one before and the other after the network formation game, in part 1 and part 3 respectively.

Each game is played with the strategy vector method (Selten, 1965). Each player is asked to make two decisions: an allocation decision, in the role of a trustor; and a return decision, in the role of a trustee. For the allocation decision in the role of a trustor, each subject is asked to allocate the entire endowment among the five other participants, each identified by the same icons used in the network formation game (see Section 3.1) and carried through the experiment's three parts. So, for each trustee in the trust-allocation game in part 3, the trustor can recall play in the network formation game and condition his decision upon such information. In this respect, icons used in the trust-allocation game in part 1 are expected to carry no meaningful information. Any amount allocated to a trustee is multiplied by three, as in Berg et al. (1995).

In each trust-allocation game, all players are endowed with either 14, 15, or 16 ECUs. More in detail, we split the 14 sessions into three different conditions and, in particular, we run 4 sessions in which trustors are endowed with 14 ECUs, 6 sessions in which they are endowed with 15 ECUs, and 4 sessions in which they are endowed with 16 ECUs. We run the three conditions to control for

⁶ Conte et al. (2015) experimentally investigate how players establish connections using the data from this network formation game also analysed in Di Cagno and Sciubba (2010).

⁷ More specifically, after the 30th round, the probability of entering a further round of play is 50%.

⁸ We identify players through seemingly unrelated icons rather than through letters or numbers in order to avoid salient connections (links in alphabetical order, for example). By avoiding labels we also prevent that considerations of homophily may drive network formation.

⁹ Screenshots of the computer screen seen during the experiment are included in the experimental instructions provided in Appendix A.

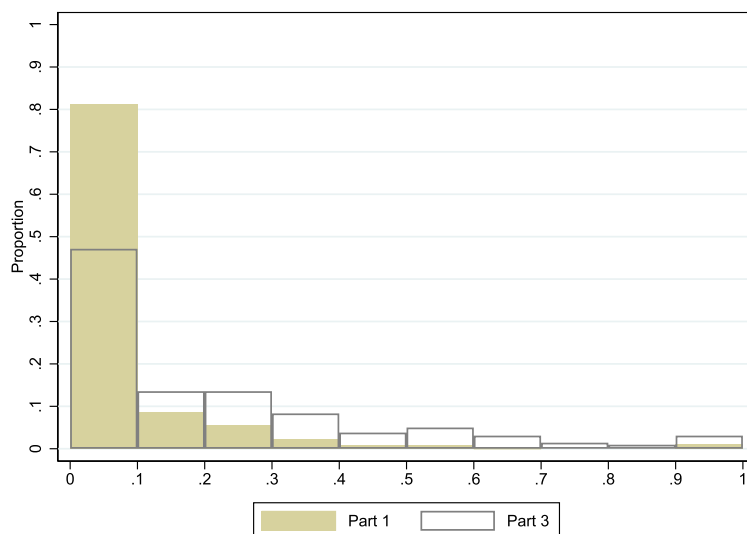


Fig. 1. Gini coefficient of trust allocation.

the fact that equality considerations or the salience of an ‘equal split’ allocation do not drive participants’ trust preferences. In the 6 sessions in which they are endowed with 15 ECUs, trustors are able to split the endowment evenly across the five trustees; in the 14- and 16-ECU conditions, trustors are forced to differentiate their offering and identify at least a trustee to receive one ECU less (in the 14-ECU condition) or one ECU more (in the 16-ECU condition) than the others.

For the return decision in the role of a trustee, each player is asked to specify how many ECUs she would return if the amount allocated would be 1 (then of 3 ECUs), 2 (6 ECUs), 3 (9 ECUs), and so on. Note that a return decision could not be conditioned upon any information about a potential trustor, neither in part 1 nor in part 3, since players are not asked to state a return contingent upon the identity of a particular player, but only upon different levels of trust.¹⁰

Given the presence of multiple trustees, a trustor earns the sum of ECUs returned by the five (at most) trustees. Trustees earn the amount received by the randomly drawn trustor, less any amount returned.

4. Data analysis and results

4.1. The allocation of trust

In the trust-allocation games, players are demanded to allocate their endowment $m \in \{14, 15, 16\}$ to the other five members of their group, $j = 1, \dots, 5$.

In order to evaluate and analyse the trust-allocation choices, for each player i , we calculate the Gini coefficient, varying from the value 0 for an equally split endowment to the value 1 when the 15-ECU endowment is allocated to one single member of the group. Instead, when the endowments are either 14 or 16 ECUs, being an equal split unfeasible, the minima of this index take values 0.071 and 0.062, respectively. Therefore, for the sake of comparison, we normalise the index, so that in the latter cases it takes values from 0 to 1 as when the endowment is 15 ECUs.

According to t tests (bootstrapped adjusting for clustering at the group level, 1000 replications), there appears to be no difference in the distributions of the Gini coefficient between endowment conditions, both in part 1 (between comparison: condition 14 vs. 15, p -value = 0.527; endowment 14 vs. 16, p -value = 0.268; endowment 15 vs. 16, p -value = 0.492) and in part 3 (between comparison: condition 14 vs. 15, p -value = 0.719; endowment 14 vs. 16, p -value = 0.320; endowment 15 vs. 16, p -value = 0.463), confirming that the allocation of trust is not dictated by the salience of a particular split. Therefore, we will conduct the following analysis without distinguishing between conditions.

Fig. 1 displays the distributions of the Gini coefficient per initial endowment in part 1 and part 3 and Table 1 reports their summary statistics. The plots are characterised by a large mass at 0 (81.19% of the observations in part 1 and 47.14% in part 3), indicating a clear preference for an even allocation of trust. The remainder allocates the endowment in different, more unequal ways, but very seldom to one player only (0.95% in part 1 and 3.10% in part 3). The distribution of the coefficient has a higher mean and median in part 3 and it is also more dispersed than that in part 1. A paired t test (bootstrapped adjusting for clustering at the group level, 1000 replications) confirms that there is a difference in the coefficient distributions between the two parts (p -value = 0.000), which is consistent with trustors directing their trust towards some trustees more than others, when the trust game is played after the network formation game.

¹⁰ It is as if for each trustee all the other group members were the trustor with 20% probability.

Table 1
Summary statistics of the Gini coefficient of trust allocation.

	observations	mean	standard deviation	median	min	max	skewness	kurtosis
Part 1	420	0.0486	0.1354	0	0	1	4.3511	26.6466
Part 3	420	0.1896	0.2505	0.1077	0	1	1.5152	4.8442

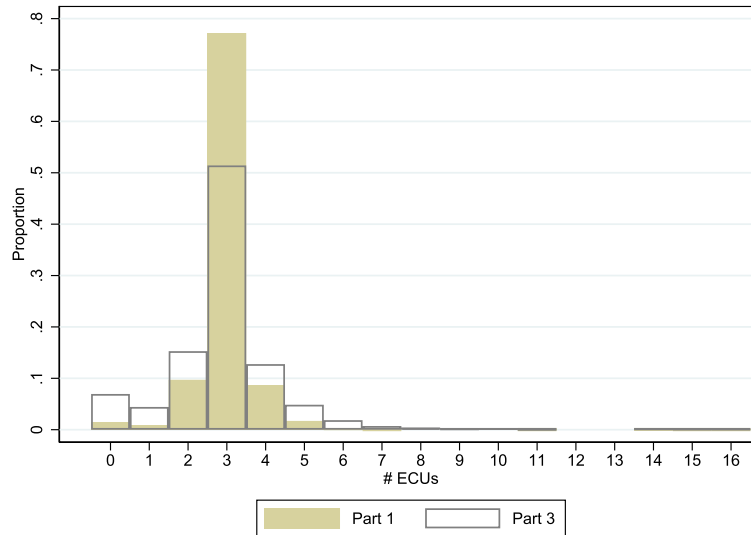


Fig. 2. ECUs allocated to each trustee.

A comparison between the trust allocations submitted in part 1 and part 3 shows that 72.14% of subjects change their strategy. This percentage remains high at 48.57%, even if we consider equivalent allocations with a simple permutation of the trustees.¹¹ Moreover, the strategies played in part 3 with respect to part 1 result in distributions concentrated on a smaller number of trustees for 61 out of 420 trustors (14.52%) and higher for only 4 (0.95%). Hence, we can conclude that trustors change their allocation of trust in part 3 compared to part 1.

Fig. 2 displays the number of ECUs which trustors allocate to trustees in part 1 and part 3 respectively. In line with what was observed in Fig. 1, we see that, in part 1, trustors predominantly (77.2%) allocate 3 ECUs to the trustees; in part 3, instead, trustors redistribute the endowment so that the number of trustees who receive a number of ECUs larger than 3 increases. Specularly, the number of those who receive less than 3 ECUs increases as well. We note that trustors in part 3 tend to distribute the endowment less equally than before, but in a way to give 4 or 5 ECUs to some, rarely more. Only in a few cases (2.5%) they assign more than 6, and in 13 cases (3.1%) the whole sum.

In summary, the results of this descriptive analysis of trust-allocation game data suggest that differences emerge when we compare how trustors allocate their trust before and after the network formation game, in that at least 50% of them appear to change their strategy, even if they rarely attach their trust to one single trustee. Whether such changes are induced by the social connections established in the course of the network formation phase is what the econometric analysis in the next two subsections is meant to shed light on.

4.1.1. The econometric model of trust allocation

Before describing the econometric model used to fit the data from choices on the allocation of trust, it is important to highlight that the numbers of ECUs from the endowment each trustor allocates to each of the trustees in his group are not independent: giving one ECU to one of the trustees prevents giving it to a different trustee. As a result, the five ‘numbers of ECUs’ allocated to the trustees are not independent and cannot be analyzed as if they were. Instead, they need to be considered as a system of non-independent equations.

For this purpose, let us try to figure out trustor i faced with the other five group members (the trustees), $j = 1, \dots, 5$, to whom he is asked to allocate 1 ECU worth of trust.

Trustor i ’s perception of j ’s trustworthiness is

$$\hat{t}_{ij} = \beta' x_{ij} + \epsilon_{ij} \tag{1}$$

¹¹ For example, let us take the five ways in which a trustor can split an endowment of 14 to the other five members in the group when he decides to entrust 3 units to four members and 2 units to one member. These five strategies are not identical, but they are equivalent as the Gini coefficient is the same.

Here, $\beta'x_{ij}$ is the deterministic component of i 's perception, which is a linear combination of the characteristics of i and j and their relationship, x_{ij} , through the vector of coefficients β (which includes a constant) representing the effects of such characteristics on i 's perception of j 's trustworthiness; ϵ_{ij} is a random component, independent and identically distributed as a Gumbel distribution (see McFadden, 1974).

Suppose that trustor i is asked to choose to whom of the five trustees he wants to allocate 1 ECU worth of trust out of his endowment m . Then, i assigns the 1 ECU to the trustee j whose perceived trustworthiness is highest. Finally, suppose that this operation is repeated m times, once for each of his units of endowment. Then, the likelihood contribution of trustor i is

$$L_i = \prod_{k=1}^m \prod_{j=1}^5 p_{ij}^{m_{ij,k}} \tag{2}$$

Here, $m_{ij,k}$ is an indicator function taking the value 1 if i assigns his 1 ECU to trustee j in trial k , 0 otherwise; $m_{ij} = \sum_{k=1}^m m_{ij,k}$ is the number of ECUs that i assigns to trustee j in m trials, where $m = \sum_{j=1}^5 m_{ij}$ is trustor i 's endowment. The probability that trustor i assigns 1 ECU worth of trust to trustee j among the five possible trustees is given by

$$p_{ij} = \frac{\exp(\beta'x_{ij})}{\sum_{j=1}^5 \exp(\beta'x_{ij})} \tag{3}$$

Guimarães et al. (2003) and Guimarães and Lindrooth (2007) show that this model is a variant of the Grouped Conditional Logit, and that by interpreting the m_{ij} 's as count variables following a Poisson distribution, it can be estimated via a simple Poisson regression. However, all the variables which are constant within the set of alternatives (the five trustees), that means i 's fixed effects and the fixed effects for the group to which i belongs cancel out, and cannot be evaluated.

The full-sample log-likelihood for the set of allocation strategies chosen by the 420 trustors in part 3 is given by:

$$\log L(\beta) = \sum_{i=1}^{420} \ln L_i \tag{4}$$

The model is estimated using the allocation strategies chosen by the trustors in part 3. Our sample consists of 420 trustors¹² and each strategy consists of a 5-vector containing the ECUs assigned to each of the five trustees, making a total of 2100 observations. To estimate the model, the method of Maximum-Likelihood is used.

4.1.2. Estimation results

Table 2 reports the estimation results of the model described in the previous section, using data collected in the trust-allocation game played in part 3 of the experiment.¹³ All the distributions of the explanatory variables are displayed in Appendix B.

The estimation results show that the coefficient on the 'proportion of rounds i and j have been directly linked' is positive and statistically significant. However, the coefficients on the 'proportion of rounds i and j have been only indirectly linked', the 'proportion of rounds j has been isolated' (i.e., ended the round with no links), and the 'average number of j 's indirect links' are negative and statistically significant. Lastly, the coefficient on the 'average number of j 's direct links' is positive but lacks statistical significance.

The estimated coefficients in Table 2 may not directly indicate the effect of a variable change on the probability that trustor i chooses to assign 1 ECU worth of trust to trustee j among the five possible trustees in their group, denoted as p_{ij} in Eq. (3). This is because they measure the effect of a change in the explanatory variables in x_{ij} (see Eq. (1)), that appears on both the numerator and the denominator of Eq. (3). Therefore, to assess the impact of social network interaction on trust, we calculate the marginal effect of a change in the explanatory variables (one at a time) on p_{ij} . In doing so, we chose to calculate the marginal effects of variables that have a more concrete meaning with respect to round proportions and averages. For example, instead of measuring the impact of a change in the proportion of rounds that i and j have been linked, we calculated the effect of i and j being directly linked in one additional round. We made this choice because a change of, for instance, 10 percentage points in terms of round proportion corresponds to a 10-round increase when the total rounds played are 100, but only a 1-round increase when the total rounds played are 10. This same rationale applies to the contribution of other variables to proportions of rounds and averages.

The marginal effects, computed from specification (all rounds) in Table 3 and expressed in percentage points, are reported in Table 3. They indicate positive and significant effects of an additional round in which i and j are linked and of a unit increase in the number of j 's direct links per round in the network formation game on the probability of i allocating trust to j .¹⁴ The marginal effects of all the other variables are negative and statistically significantly different from 0. The magnitude of the marginal effects

¹² Remember that each individual in our sample plays both as a trustor and as a trustee, using the strategy-vector method. Therefore, there are a total of 420 trustors and an equal number of trustees.

¹³ In Table 2, we report the specification that, in our opinion, best summarises the estimation results. Additional specifications (which also include regressors calculated on the last few rounds of the network formation game to consider the possibility that subjects may base their allocation decisions in part 3 on their more recent history in that game) are available upon request from the authors.

¹⁴ We note that the significant effect of a unit increase in the number of j 's direct links has a statistically significant, positive effect on the probability of i assigning an ECU worth of trust to j *per se*. However, being such an effect very small and so the contribution of a unit increase to the average number of direct links, the total effect of a change in the latter on j 's valuation is small and not statistically significant, in line with Table 2.

Table 2
Estimation results of network interaction on trust.

Dependent variable: vector of ECUs assigned to other group members in the trust-allocation game	
<i>Explanatory variables:</i>	
Proportion of rounds <i>i</i> and <i>j</i> have been directly linked	0.3492*** (0.0951)
Proportion of rounds <i>i</i> and <i>j</i> have been only indirectly linked	-0.2722* (0.1250)
Proportion of rounds <i>j</i> has been isolated	-0.5880** (0.2045)
Average number of <i>j</i> 's direct links	0.0406 (0.0377)
Average number of <i>j</i> 's indirect links	-0.0750* (0.0311)
Constant	1.2316*** (0.0911)
# subjects	420
# observations	2100
Log-likelihood	-3772.9913

Note: The marginal effects are based on the estimation results in Table 2, and expressed in percentage points.

* *p*-value < 0.05; ** *p*-value < 0.01; *** *p*-value < 0.001.

Standard errors are clustered at the individual level.

Table 3
Marginal effects of network-related variables on trust-allocation probability.

<i>i</i> and <i>j</i> have been directly linked in one additional round	0.1760*** (0.0170)
<i>i</i> and <i>j</i> have been only indirectly linked in one additional round	-0.1380*** (0.0115)
<i>j</i> has been isolated in one additional round	-0.3005*** (0.0094)
The number of <i>j</i> 's direct links increases by one	0.0209*** (0.0007)
The number of <i>j</i> 's indirect links increases by one	-0.0385*** (0.0012)

Note: The marginal effects are based on the estimation results in Table 2, and expressed in percentage points.

* *p*-value < 0.05; ** *p*-value < 0.01; *** *p*-value < 0.001.

Standard errors are clustered at the individual level.

is mostly small. This is consistent with what was discussed in Section 4.1 regarding the not overwhelming but clear, relevant and statistically significant reallocation of trust after the network formation game.

Result 1: *A direct, long-lasting relationship between the trustor and the trustee directs trust towards the trustee. Trustees with a large number of connections attract trust but only marginally. In contrast, indirect relationships, and particularly isolation, direct trust away from the trustee.*

4.2. Trustworthiness and returns to trust

We now look at responses to trust by trustees in the trust-allocation games. We find that trustworthiness remains stable between the two trust-allocation games. Fig. 3 displays the distribution of individual-wise correlations between the ECUs returned in parts 1 and 3 per amount received. We observe that for over 75% of the subjects, this correlation exceeds 0.9, and for more than 90% of the subjects, it is higher than 0.75.

When we compare amounts returned in parts 1 and 3, we find that trustees tend to return less in part 3 than in part 1. This is shown in Fig. 4 which displays box plots per ECUs received (i.e. ECUs allocated multiplied by three) in part 1 and part 3 (in blue and red, respectively). This is also confirmed from paired *t* tests per received ECUs. All the tests of the null hypothesis that the ECUs returned in part 1 and part 3 are equal against the alternative that they do not have a *p*-value < 0.001, except for when the received ECUs are 3 and 48 with *p*-values = 0.1203 and = 0.0251. On average, in part 1, trustees return to the trustor 20.25% (stand. err. 0.23) of the received ECUs and 17.09% (stand. err. 0.21) in part 3, which are significantly different (paired *t* test: *p*-value = 0.000).

This result stems from two distinct processes: i) On one side, the probability of the trustee returning trust—giving back a positive amount of ECUs in response to what was received from the trustor—increases with the number of ECUs received (refer to Fig. 5). Notably, this probability tends to be lower in part 3 compared to part 1. ii) On the other side, whenever the trustee returns trust, she tends to return a proportion of the ECUs received that is independent of the amount received, always lower in part 3 than in part 2, on average (see Fig. 6).

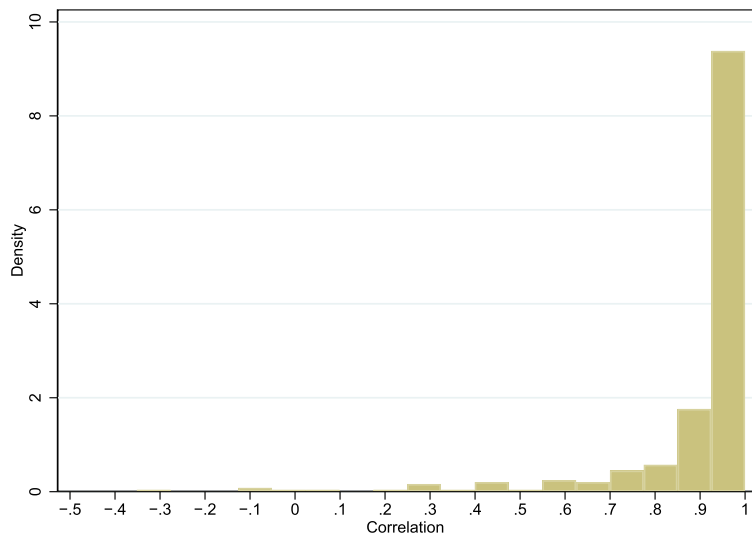


Fig. 3. Histogram of the correlation between ECUs returned in Part 1 and ECUs returned in Part 3, by individual.

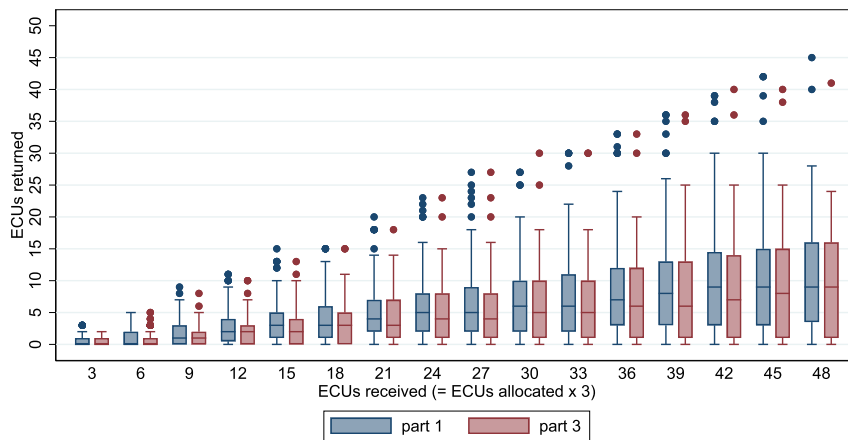


Fig. 4. Box plots of the ECUs returned to the trustor by the trustee in the trust-allocation game against ECUs received. *Note:* The amount received, displayed on the horizontal axis, is the amount allocated by the trustor multiplied by 3. Dots are outside values. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

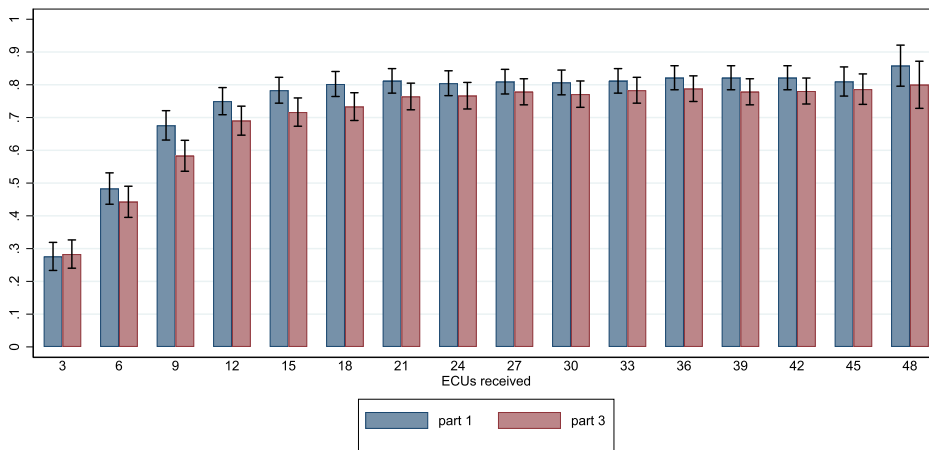


Fig. 5. Probability of being trustworthy. *Note:* The amount received, displayed on the horizontal axis, is the amount allocated by the trustor multiplied by 3. Capped spikes represent 95% confidence intervals.

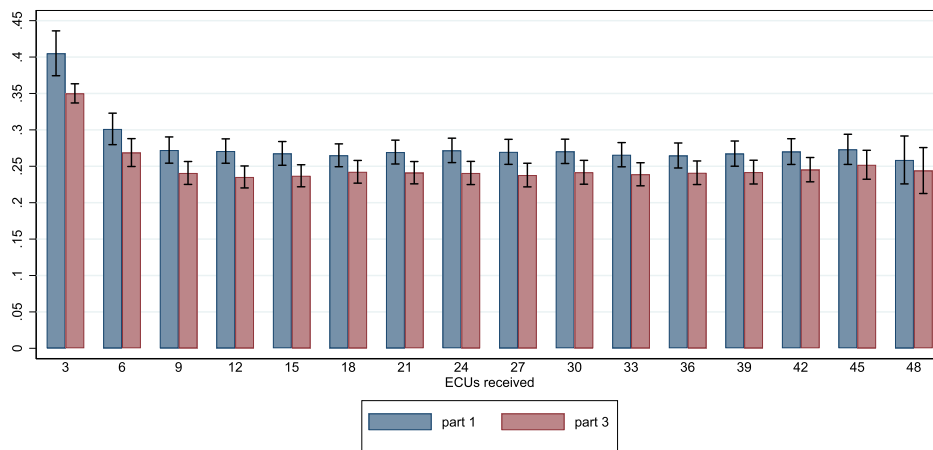


Fig. 6. Proportion of ECUs returned by trustworthy recipients. Note: The amount received, displayed on the horizontal axis, is the amount allocated by the trustor multiplied by 3. Capped spikes represent 95% confidence intervals.

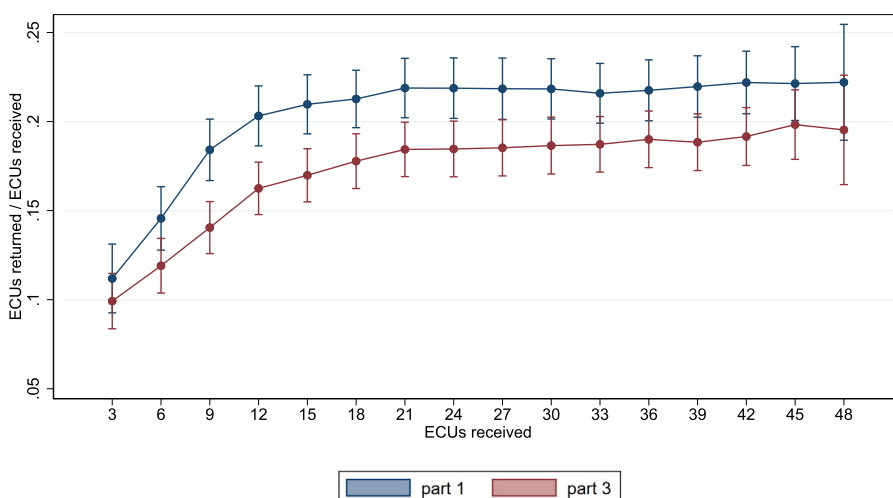


Fig. 7. Mean ratio ECUs returned / ECUs received per ECU received in the trust-allocation game. Note: Connected dots represent averages. Capped spikes represent 95% confidence intervals.

Fig. 7 shows the proportion of ECUs returned per ECU received against ECUs received for both part 1 and part 3 of the trust-allocation game. We notice that, in both parts, the proportion of ECUs returned to the trustor is larger, on average, the larger the amount of ECUs received: trustees respond more generously to generous offers. However, the extent to which they do so is lower when the trust-allocation game is played after the network formation game than when it is played before.

Result 2: On average, trustworthiness in part 3 is highly correlated to trustworthiness in part 1. However, average trustworthiness is significantly lower after the network formation game.

Given Result 2, it is not surprising that returns to trust, or trustors’ earnings, are lower in part 3 than in part 1. Trustors earn 8.424 ECUs (stand. dev. 3.396) in part 1 and 6.419 ECUs (stand. dev. 3.309) in part 3. Their distributions significantly differ according to a paired *t* test (*p*-value = 0.000). In 310 cases (73.8%), trustors earn less in part 3, and in only 56 cases (13.3%) they earn more.

Considering trustees’ responses to trust, we note that the optimal strategy for trustors, aiming to maximise earnings, is to concentrate trust on a smaller number of trustees. Fig. 7 appears to suggest that concentrating trust is even more beneficial in part 3 than in part 1, since the distribution of the ratio of ECU received and returned is strictly increasing in part 3, while it peaks at 24 and then remains constant in part 1. However, such a difference is not statistically significant. Hence the lower returns to trust in part 3 compared to part 1 cannot be ascribed to trustors acting less strategically or identifying the optimal (selfish) trust allocation strategy any less in part 3 than they do in part 1.

The diminished returns to trust in part 3 are unrelated to learning effects, as our design precludes them by ensuring that subjects do not receive any feedback on the game played in part 1 until the end of the experimental session. In particular, subjects are not told which role they will play and the related payment and do not know the strategy choices made by the other participants in their

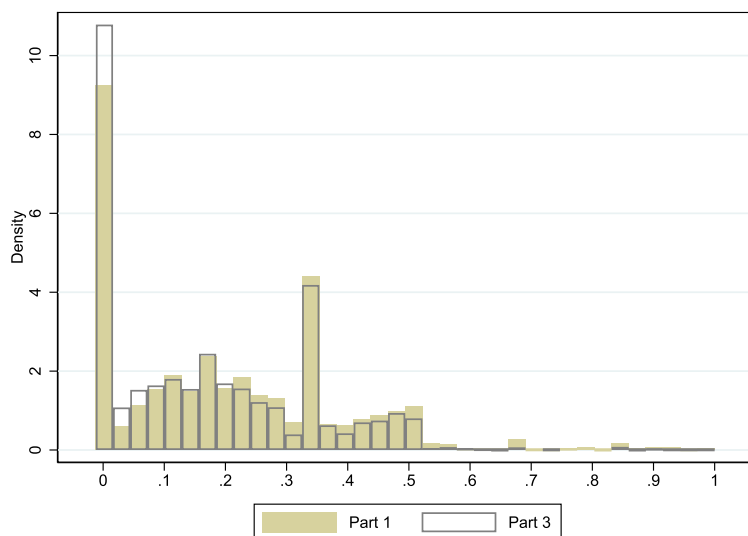


Fig. 8. Histograms of the ratio ECUs returned / ECUs received.

group. Moreover, they do not know whether part 1 or part 3 is relevant for payment until the end of the experiment. Hence, they have no way of evaluating how their strategy choice in part 1 has performed, which rules out that learning effects take place between part 1 and part 3.

A possible explanation for the lower returns to trust in part 3 is the income effect induced by the network formation game in part 2, where all participants earn a positive amount of ECUs. Instead, when the trust-allocation game is played first, subjects do not know the characteristics of the games that they are going to face afterwards, and cannot anticipate such an effect. We mitigate income effects between part 2 and part 3 with the fact that only one round (out of the 30 or more played) is selected for payment and therefore participants remain very uncertain about their part 2 earnings until the very end of the experimental session.

Another possible explanation for the lower returns to trust in part 3 compared to part 1 could be the expectations of trust generated by the network formation stage. For example, receiving 9 ECUs in a completely anonymous setting may be perceived as more (or less) trusting compared to receiving the same amount of ECUs in part 3, when the network formation stage has been characterised by strong (or weak) connections. Although the strategy method reduces the potential for relational effects, it does not rule them out completely.

However, in order to clean the results from any remaining income and/or expectations effects, we calculate the earnings for each trustor from the strategies submitted in parts 1 and 3 based on trustees' trustworthiness in part 1. In other words, we calculate the earnings they achieved before the network formation game and those they would have achieved before the network formation game, had they adopted the same strategy used in the trust-allocation game in part 3.

The mean earnings from the choices made in part 1 were 8.424 ECUs (stand. dev. 3.396); those made after would have been worth 8.295 ECUs (stand. dev. 3.896). The two distributions are not significantly different (paired t test, clustered at group level: $t = 1.040$, p -value = 0.315).¹⁵ This comparison across earnings suggests that trustors derive no monetary advantage from directing trust towards those trustees that they perceive as more trustworthy.

4.2.1. The econometric model of trustworthy behaviour

We estimate the effects of a trustee's characteristics as observed in the network formation game on her trustworthiness, measured as the ratio of ECUs returned to ECUs received. The histograms of these measures are displayed in Fig. 8. The plot shows remarkable similarity between part 1 and part 3, except for a higher proportion of 0 returns in the latter (30.32%) compared to the former (25.97%). In both parts, this proportion decreases as the number of ECUs received increases, consistent with Fig. 7, suggesting that trustees may prefer to reach a minimum threshold of ECUs before initiating trust returns.

To account for this data characteristic, we employ a zero-inflated model, which separates the estimation of the proportion of ECUs returned per ECU received from the factors influencing the decision to offer something positive. Failing to do so might lead to an erroneous amplification of the effects of variables that impact both the decision to offer something positive and the proportion of ECUs offered in the same direction, positively or negatively. Conversely, variables with opposing impacts on these two aspects of the decision may offset each other.¹⁶

¹⁵ For the most curious readers, on the basis of the ECUs received in part 3 of the experiment, the mean earnings of the trustor from the strategies played before the network formation game were 6.517 ECUs (stand. dev. 3.031), while those played after were 6.419 ECUs (stand. dev. 3.309). The two distributions do not show a significant difference (paired t test, clustered at the group level: $t = 0.907$ and p -value = 0.423).

¹⁶ This concept is extensively explained in the literature on endogenous sample selection, with references to, among others, Heckman (1979).

Specifically, the model employed here is a zero-inflated log-linear regression, controlling for individual fixed effects through individual-specific regressors observed in the network formation game.¹⁷ The underlying concept of this model is that, when a trustee receives a certain number of ECUs, she decides whether to retain everything (because the amount does not reach the minimum threshold) or to respond to the allocated trust by returning part of the ECUs received to the trustor.

Let $t_{j,k}$ denote the proportion of ECUs returned by trustee j with respect to ECUs received from the trustor and $\Pr(t_{j,k} = 0)$ the probability of j returning 0 ECUs (that is $t_{j,k} = 0$), $\Pr(t_{j,k} > 0)$ being trustee j 's probability of returning trust ($t_{j,k} > 0$), when the ECUs received are $k \in \{3, 6, \dots, 3m\}$, with $m \in \{14, 15, 16\}$ for the three endowment conditions. Essentially, while $\Pr(t_{j,k} > 0)$ captures trustee j 's probability of being trustworthy (Fig. 5), $t_{j,k}|t_{j,k} > 0$ measures the extent of the trustee's trustworthiness (see Fig. 6).

$\Pr(t_{j,k} > 0)$ is modelled via a normal probability distribution function, as in a standard probit model, and the level of trust returns as a proportion of the trust received ($t_{j,k}|t_{j,k} > 0$) is driven by a lognormal distribution. Both of them depend on the number of ECUs received and the subject's characteristics measured in the network formation game ($r'_{j,k}$ and $z'_{j,k}$, respectively), via the vectors of coefficients δ and θ , respectively. The individual likelihood contribution is then given by

$$L_{j,k} = \Phi(-r'_{j,k}\delta)^{I_{(t_{j,k}=0)}} \times \left[\Phi(r'_{j,k}\delta) \times \frac{1}{\sigma t_{j,k}} \phi\left(\frac{\log(t_{j,k}) - z'_{j,k}\theta}{\sigma}\right) \right]^{I_{(t_{j,k}>0)}} \tag{5}$$

Here, $\phi(\cdot)$ and $\Phi(\cdot)$ represent the standard normal density and the distribution functions, respectively; I is an indicator function taking the value 1 if the condition specified in the pedix is true, and 0 otherwise.

The full-sample log-likelihood for the vector of 'ECUs returned over ECUs received' per ECU received by the 420 trustees in part 3 is given by

$$\log L(\beta) = \sum_{j=1}^{420} \ln \left[\prod_k L_{j,k} \right] \tag{6}$$

The model is estimated using the data collected via the strategy vector method from subjects in the role of trustees as a response to possible allocated trust in part 3. Our sample consists of 420 trustees each submitting an m -vector of responses to allocated trust, making a total of 6300 observations. To estimate the model, the method of Maximum-Likelihood is used.

4.2.2. Estimation results

Table 4 displays the estimation results of the model described in Section 4.2.1. Results are organised in two columns: the first reports coefficient estimates for the probability of the trustee returning trust, $\Pr(t_{j,k} > 0)$; the second refers to the extent of the trustee's trustworthiness, $\log(t_{j,k})|t_{j,k} > 0$, respectively. The trustee's fixed effects are obtained from her history in the network formation game.

The reported specifications include the same explanatory variables as in Table 2, excluding those related to i , the trustor, as they are unknown to the trustee. We did this to investigate whether the characteristics of the trustee that attract more trust also account for trustworthiness,¹⁸ and also condition on the number of ECUs received and its transformation.

Overall, our model estimates that the probability of trustworthiness increases with the logarithm of the number of ECUs received, while the proportion returned of the received ECUs decreases initially with the amount received from the trustor, and becomes constant afterwards. These two opposite effects are in line with what is observed in Figs. 5 and 6.

None of the other trustee's characteristics, perceived as signals of trustworthiness by the trustor (see Tables 2 and 3), seems to be of any relevance either on the probability of the trustee being trustworthy or on the extent of her trustworthiness, if any.

These findings demonstrate that the signals extracted by the trustor from the network formation game do not accurately reflect the individual trustworthiness of the trustee.

Result 3: *The trust-allocation strategies, conditional on signals observed during the network formation game, do not yield higher returns when no relationship has been established between trustor and trustee. Social interaction does not provide strong informative insights into individual trustworthiness.*

Table C.1 presents a different specification of the model of trustworthy behaviour. In addition to '# of ECUs received' and its transformations, it includes the dummy variable 'j performed worse than average'. This variable takes the value 1 if the trustee performed, in terms of profits, worse than the average in her group during the network formation game, and 0 otherwise. The inclusion of this specification aims to test whether the observed lower trustworthiness in part 3, compared to part 1, is due to residual income effects generated in the network formation game. If true, we would expect those who earned less in the network formation game to set their minimum threshold of ECUs before initiating trust returns higher and/or return a lower proportion of the ECUs received. This would allow them to compensate for the lower income in the network formation game. In both parts of

¹⁷ Zero-inflated models are more usually estimated from count data (Lambert, 1992). Here, we derive a version for continuous data with point mass one at 0.

¹⁸ This may not necessarily be the case, as factors other than these trustee characteristics may offer a better explanation for trustworthiness. We explored alternative variables but did not find any that performed statistically better than the ones reported in explaining trustworthy behaviour. All the results are available from the authors upon request.

Table 4
 Estimation results of network-induced effects on trustworthy behaviour.

	Pr ($t_{j,k} > 0$)	$\log(t_{j,k}) t_{j,k} > 0$
log (# of ECUs received)	0.5014*** (0.0319)	
(# of ECUs received)		-0.0668*** (0.0074)
(# of ECUs received) ²		0.0022*** (0.0003)
(# of ECUs received) ³		-0.00002*** (4.54e-06)
Proportion of rounds j has been isolated	-1.1506 (0.6856)	-0.1785 (0.4160)
Average number of j 's direct links	-0.0794 (0.1490)	-0.0592 (0.0946)
Average number of j 's indirect links	-0.0992 (0.1165)	-0.0019 (0.0727)
Constant	-0.4839 (0.5395)	-0.9207** (0.3385)
σ		0.5053*** (0.0163)
# observations	6300	
# subjects	420	
Log-likelihood	-6009.5421	

Note: The dependent variable is the logarithm of the ratio of ECUs returned by the trustee to ECUs received from the trustor in the trust-allocation game in part 3. Standard errors (in parentheses) are clustered at the individual level.
 * p -value < 0.05; ** p -value < 0.01; *** p -value < 0.001.

the model, the coefficient on the aforementioned variable is not significantly different from 0. Therefore, we can exclude what was observed as the result of an income effect, giving more strength to the expectation of greater trust after the network formation game as an explanation of the lower trustworthiness observed.

5. Conclusion

In this paper, we have explored experimentally the individual and relational components of trust and trustworthiness in a network formation game where members of a group can establish (either directly or indirectly) mutually beneficial connections with one another.

We measure trust through an involuntary trust game *à la* McCabe et al. (2003) in which trustors have to allocate their given endowment to several trustees, with no option of retaining the remainder for themselves. We compare allocations of trust before and after the formation of a social network finding that, as we expected (Hypothesis 1), trustors significantly change their allocation of trust after experiencing the network formation game by concentrating it towards those trustees that they perceive as more trustworthy. This could be either because trustors infer signals of individual propensity to being trustworthy from trustees' behaviour during the network formation game, or because they wish to reciprocate positive behaviours with greater trust. To sharpen our predictions, we develop a theoretical model of trust and trustworthiness which allows for both individual characteristics and reciprocal behaviour and we illustrate how both motives may be driving trustors' allocations concurrently.

In our empirical model of trust allocation, we confirm that the probability of allocating trust to a particular trustee increases with the number of rounds in which trustor and trustee have been directly linked and decreases with the number of rounds in which trustor and trustee have been linked indirectly and the average number of indirect links enjoyed by the trustee. Interestingly, the number of direct links established by the trustee is not significant, but the allocation of trust significantly decreases with the number of rounds in which the trustee has been isolated. Trustees who have been able to establish stable direct relationships with the trustor are perceived as more trustworthy than those who have been free riding on indirect connections or have failed to network altogether.

We measure trustworthiness by eliciting trustees' responses to trust using the strategy vector method in which we condition trustee responses to the number of ECUs received and not also to the identity of the trustors. This allows us to focus on individual trustworthiness alone, abstracting from any desire that trustees may have to reciprocate with more (or less) trustworthiness positive (or negative) behaviour that they might have witnessed (from a particular trustor) in the network formation game.

This design allows us to elicit the individual component of trustworthiness which we expect will remain unaffected by the network formation game (Hypothesis 2).

We find a strong individual-wise correlation of trustworthiness before and after the network formation game, but the average level of trustworthiness is significantly lower when the trust-allocation game is played after the network formation game. Our design excludes that this may be due to learning effects (feedback is only given to subjects at the end) and mitigates the possibility that it may be due to income effects (only one of the two trust games and one of the rounds of the network formation game are randomly selected for payment). We argue that a residual explanation for the reduced trustworthiness in part 3 is related to an expectation

effect: trustees anticipate that trustors will concentrate trust on a smaller number of recipients after network formation and respond less positively to amounts of trust that would have triggered a trustworthy response in an anonymous setting. We decompose our measurement of trustworthiness by looking at the probability of returning a strictly positive amount and the share of amount received which is returned. We find that both are significantly lower after network formation: trustees' positive responses are triggered by higher levels of trust; also, the extent of positive responses is lower. Both effects are compatible with our interpretation of expectation of greater trust after network formation.

Our empirical model of trustworthiness confirms that network variables are not significant and that the only determinant of trustworthiness both before and after the network formation game is the amount received.

Our third hypothesis, and corresponding result, focuses on trustworthiness from the point of view of the trustors, i.e., on returns to trust. If indeed trustors were able to infer from trustees' behaviour during the network formation game informative signals on their individual trustworthiness, we would expect returns to trust to be higher after network formation (Hypothesis 3).

Given that we find that average trustworthiness is lower after network formation, it is not surprising that trustors' earnings decrease too. To purge our results of the expectation effect discussed above, and of any residual learning or income effects, we compare trustors' earnings in part 1 to what trustors' earnings would have been if trustors had been able to allocate trust as they then did in part 3, with trustees responding as they did in part 1.

We find that there is not a significant difference in trustors' earnings. This implies that the signals of perceived trustworthiness inferred during the network formation game are not informative about individual levels of trustworthiness: by engaging with their potential trustees in network formation, trustors are unable to gauge who is more likely to return positive and larger amounts with any more precision than when they allocate in absence of information.

This result is only in apparent contrast with what has been previously emphasised in the literature on trustworthiness and social networks. Di Cagno and Sciubba (2010) show that returns to trust are increased when a trust game is played after a network formation game. Field experiments¹⁹ have documented how trust and, to an even greater extent, trustworthiness are stronger when trust games are played by subjects who are closer in a social network of ties. Our results can be interpreted as evidence that the increased returns to trust observed among agents which are close in a social network, do not depend on reduced informational asymmetries, and on trust being better allocated in a social network, but rather on the fact that trustworthiness is mainly a relational characteristic and that the social network channels reciprocal behaviour.

We may think we have chosen our friends so that they are kind, whereas it is their friendship that drives kind reciprocation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

We thank the editor, Professor Daniela Puzzello, and two anonymous referees. Assistance for writing the experimental software and conducting the experiment from Albrecht Noll is gratefully acknowledged, as well as financial support from the Max Planck Institute of Economics, Jena, Germany. The usual disclaimer applies.

Appendix. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jebo.2023.12.031>.

References

- Ashraf, N., Bohnet, I., Piankov, N., 2006. Decomposing trust and trustworthiness. *Exp. Econ.* 9 (3), 193–208.
- Barr, A., 2003. Trust and expected trustworthiness: experimental evidence from zimbabwean villages. *Econ. J.* 113 (489), 614–630.
- Berg, J., Dickhaut, J., McCabe, K., 1995. Trust, reciprocity, and social history. *Games Econ. Behav.* 10 (1), 122–142.
- Bicchieri, C., Xiao, E., Muldoon, R., 2011. Trustworthiness is a social norm, but trusting is not. *Polit. Philos. Econ.* 10, 170–187.
- Binzel, C., Fehr, D., 2013. Social distance and trust: experimental evidence from a slum in Cairo. *J. Dev. Econ.* 103, 99–106.
- Camerer, C.F., 2003. *Behavioral Game Theory*. Princeton University Press.
- Conte, A., Di Cagno, D.T., Sciubba, E., 2015. Behavioral patterns in social networks. *Econ. Inq.* 53 (2), 1331–1349.
- Di Cagno, D., Sciubba, E., 2010. Trust, trustworthiness and social networks: playing a trust game when networks are formed in the lab. *J. Econ. Behav. Organ.* 75 (2), 156–167.
- Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10 (2), 171–178.
- Glaeser, E.L., Laibson, D.I., Scheinkman, J.A., Soutter, C.L., 2000. Measuring trust. *Q. J. Econ.* 115 (3), 811–846.

¹⁹ See, among others, Glaeser et al. (2000), Leider et al. (2009) and Binzel and Fehr (2013).

- Granovetter, M., 1985. Economic action and social structure: the problem of embeddedness. *Am. J. Sociol.* 91 (3), 481–510.
- Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. *J. Econ. Sci. Assoc.* 1 (1), 114–125.
- Guimarães, P., Figueirdo, O., Woodward, D., 2003. A tractable approach to the firm location decision problem. *Rev. Econ. Stat.* 85 (1), 201–204.
- Guimarães, P., Lindrooth, R.C., 2007. Controlling for overdispersion in grouped conditional logit models: a computationally simple application of Dirichlet-multinomial regression. *Econom. J.* 10 (2), 439–452.
- Heckman, J.J., 1979. Sample selection bias as a specification error. *Econometrica* 47 (1), 153–161.
- Jacobs, J., 1961. *The Death and Life of Great American Cities*. Random House, New York.
- Lambert, D., 1992. Zero-inflated Poisson regression, with an application to defects in manufacturing. *Technometrics* 34 (1), 1–14.
- Leider, S., Möbius, M.M., Rosenblat, T., Do, Q.-A., 2009. Directed altruism and enforced reciprocity in social networks. *Q. J. Econ.* 124 (4), 1815–1851.
- McCabe, K.A., Rigdon, M.L., Smith, V.L., 2003. Positive reciprocity and intentions in trust games. *J. Econ. Behav. Organ.* 52 (2), 267–275.
- McFadden, D., 1974. Conditional logit analysis of qualitative choice behavior. In: *Frontiers in Econometrics*. Academic Press, New York, pp. 105–142. Chapter 4.
- Ostrom, E., Walker, J., 2003. *Trust and Reciprocity: Interdisciplinary Lessons for Experimental Research*. Russell Sage Foundation.
- Putnam, R.D., 2000. *Bowling Alone: The Collapse and Revival of American Community*. Simon et Schuster, New York.
- Riyanto, Y.E., Jonathan, Y.X., 2018. Directed trust and trustworthiness in a social network: an experimental investigation. *J. Econ. Behav. Organ.* 151, 234–253.
- Salgado, M., Núñez, J., Mackenna, B., 2021. Expectations of trustworthiness in cross-status interactions. *Soc. Sci. Res.* 99, 102596.
- Selten, R., 1965. Die strategiemethode zur erforschung des eingeschränkt rationalen verhaltens im rahmen eines oligopolexperimentes. *Seminar für Mathemat. Wirtschaftsforschung u. Ökonometrie*.
- Torbiörn, I., 2002. Olikhet som möjlighet. *Ledmotiv* 1, 59–70.
- Tullberg, J., 2008. Trust—The importance of trustfulness versus trustworthiness. *J. Socio-Econ.* 37 (5), 2059–2071.