JENA ECONOMIC RESEARCH PAPERS

\# 2013 – 044

You Are Who Your Friends Are: An Experiment on Trust and Homophily in Friendship Networks

by

Fabian Winter
Mitesh Kataria

www.jenecon.de

ISSN 1864-7057

The JENA ECONOMIC RESEARCH PAPERS is a joint publication of the Friedrich Schiller University and the Max Planck Institute of Economics, Jena, Germany. For editorial correspondence please contact markus.pasche@uni-jena.de.

Impressum:

Friedrich Schiller University Jena
Carl-Zeiss-Str. 3
D-07743 Jena
www.uni-jena.de

Max Planck Institute of Economics
Kahlaische Str. 10
D-07745 Jena
www.econ.mpg.de

You Are Who Your Friends Are: An Experiment on Trust and Homophily in Friendship Networks

Fabian Winter* and Mitesh Kataria†‡

October 22, 2013

Abstract

We study the existence of homophily (i.e. the tendency for people to make friends with people who are similar to themselves) with respect to trustworthiness. We ask whether two friends show similarly trustworthy behavior towards strangers, and whether this is anticipated by outsiders. We develop a simple model of bayesian learning in trust games and test the derived hypotheses in a controlled laboratory environment. In the experiment, two trustees sequentially play a trust game with the same trustor, where the trustees depending on treatmen are either friends or strangers to each other. We affirm the existence of homophily with respect to trustworthiness. Trustors' beliefs about the trustees' trustfulness are not affected by the knowledge about the (non-)existent friendship between the trustees. Behaviorally, however, they indirectly reciprocate the (un-)trustworthy behavior of one trustee towards his/her friends in later interactions.

Keywords: social networks homophily trust friendship indirect tit-for-tat

JEL Classification C92 · D83 · J24 · J40

*Max-Planck Institute of Economics, Kahlaische Strasse 10, 07743 Jena, Germany. winter@econ.mpg.de.

†Max-Planck Institute of Economics, Kahlaische Strasse 10, 07743 Jena, Germany. kataria@econ.mpg.de.

‡Both authors contributed equally to this work and jointly share first authorship. We thank Vittoria Maria Levati and Michael Trost for valuable comments to the manuscript, and Ria Stangneth, Lisa Schöttl, Dvin Galstian Pour and Andreas Matzke for valuable research assistance. No IRB approval is required at the authors' institution. The usual disclaimer applies.

Society is not composed of culturally alienated beings. In dealing with someone you learn something not only about him, but also about others in his society. You learn something about population statistics.

Dasgupta (1988)

With the increasing awareness about the role of social capital in economics there has been a great interest in understanding trust. Incomplete information and contracts are a key feature in most economic transactions, which creates transaction costs. Trust thus becomes a necessary condition for economic activities to take place in such situations (Arrow, 1974). Economists traditionally talk about trust either at the micro level of individuals or the macro level of societies. In this paper we will look at the trust problem from a social network perspective which relates the micro and the macro perspective: individual decision making is embedded in macro structures of networks. Despite the widespread realization of the importance of networks in sociology (see for instance the seminal paper on job recruitment by Granovetter (1973)), networks have been largely overlooked by economist until the last decade (Jackson and Watts, 2002; Jackson and Rogers, 2007). A recent interest in economics has, however, manifested itself in several contributions investigating the relationship between social networks and other-regarding preferences (Leider et al., 2009, 2010; Goeree et al., 2010). We extend this economic literature by investigating the composition of networks and their consequences for trust relations. Assessing the quality of a job candidate or the reliability of business partners are only two economic activities where trust plays a vital role and can be affected by agents that are embedded in social networks (Kataria and Winter, 2013). We conjecture that crucial information about opinions as well as behavior of its members travel through social networks, and thus resolve some of the uncertainty about the individuals involved.

Following Coleman (1990), we will talk about trust as a two dimensional concept: A *trustor* can be *trustful*, that is expect that the person she trusts (the *trustee*) will return the trust. A *trustee* can either prove to be *trustworthy* by returning the trust granted by the trustor, or not. These trust relations are often embedded in social networks, a social structure made up of individuals or organizations (*nodes)*, connected by one or more specific types of interdependencies (*ties*), such as friendship, exchange of information or economic transactions. For simplicity, our study focuses on three connected individuals, two trustees and one trustor, which is often considered the minimal foundation of a network. Figure 1 depicts the basic network ties we investigate: The trustor (top node, denoted λ) is linked to two trustees (λ_1 and λ_2 respectively) via a mutual trust relationship. Depending on the experimental conditions, λ_1 and λ_2 may be directly connected via two sorts of links: The may (*i*) be real-life friends with each other and thus share an undirected *friendship link*, and (*ii*) λ_2 may learn about the choice of λ_1 in the experiment and thus share a directed *information link* with λ_1.

We contribute to the literature by experimentally examining how friendship affects the two dimensions of trust. Our experimental study answers three questions: (i) Is a trustee's trustworthiness towards a trustor affected if the trustee observes the (un-)trustworthy behavior of his/her friend towards this trustor? As a control, we compare if a trustee's trustworthiness is affected less if the observed person is *not* a friend of

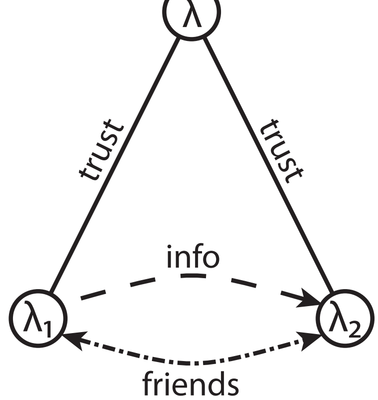

Figure 1 Basic network relations in this study. The trustor λ may enter a trust relationship, first to trustee λ_1, then to trustee λ_2, while λ_1 and λ_2 may have informational and friendship ties.

the trustee. (ii) In a second step, we extend the question as to whether trustworthiness is shared among friends even if they are ignorant about their friend's behavior. This question can also be rephrased as to whether homophily with respect to trustworthiness exists unconditional of shared knowledge. (iii) Finally, we ask whether the information about the trustees' (non-)existing friendship link is utilized by the trustors.

The rest of the paper is structured as follows: Section 1 introduces the concept of homophily in networks and corresponding hypotheses about the trustees' behavior, section 2 specifies a simple model of bayesian learning in trust games and derives hypotheses about the trustors' behavior, section 3 describes the experimental design and treatments, section 4 reports the results and section 5 concludes.

1 Homophily in trustworthiness

Homophily can be captured by the old saying that "birds of a feather flock together". It refers to the frequent finding (Bott, 1928; Loomis, 1946) that agents in a network have a higher-than-random chance of being linked to someone similar to themselves (for a review see McPherson, Smith-Lovin and Cook, 2001). Homophily may be the result of either *selection* – we make friends with people like us – or *adaptation*, that is, once we are acquainted to someone, we tend to adapt his or her behavior, manners or opinions (see Baccara and Yariv, 2013, for a theoretical model). It may be based on shared values such as political attitudes or aspirations for higher education, but also on acquired status, e.g. religion, education, occupation, or on ascribed status like race, ethnicity, age, or gender (Lazarsfeld and Merton, 1954).

Stable trust relationships are an important *outcome* of homophily: People tend to trust similar others, for instance in when asking for help in emergencyor loaning money

(McPherson, Smith-Lovin and Cook, 2001), when discussing important personal matters (DiPrete et al., 2011), or even trusting product rating sites on the internet (Golbeck, 2009; Tang et al., 2013). Not surprisingly, mutual trust is a key foundation of close friendships (Hatfield, 1984).

Extending the existing literature, we are interested in trust as a source of homophily. Igarashi et al. (2008) find in a survey study that embeddedness in closed networks, where friends of friends are usually also friends with each other, is positively correlated with *trustfulness*. Yet, little is known about the connection between network composition and *trustworthiness*. This study is a first emptemt to answer the question as to whether trustworthiness is shared among friends in a network.

We conceptually distinguish between two kind of homophily with respect to the trustees' behavior, and introduce corresponding hypotheses:

Hypothesis 1 (Conditional Homophily) *Observing a friends' behavior has a stronger effect on a trustee's choice than observing a random stranger's choice.*

In a second step, we are interested in an even stronger case, namely whether we can also find evidence that homophily is present even if people are not informed about their friends behavior. If trustworthiness is an organizing feature of networks, we could expect a positive correlation in the trustworthiness of friends, even if the friends are ignorant of their respective counterpart's choices. Consequently, our second hypothesis states that

Hypothesis 2 (Unconditional Homophily) *There is a higher correlation in choices among friends than among strangers, even in the absence of information about some other person's behavior.*

Our research hypothesis is that trustworthiness is shared among friends and we expect conditional homophily to be stronger than unconditional homophily.

2 A model of bayesian learning in trust games

In the simplest case of a trust relation there is only one type of trustees and trustors and it can be modeled as a trust game (Dasgupta, 1988, see Figure 2a). The unique equilibrium of the game is determined by backward induction: Because the trustee prefers not to return trust ($1 > r$), the trustor prefers not to grant trust in the first place ($s > 0$), despite the fact that this equilibrium yields a lower payoff than the mutually cooperative outcome of granting and returning trust. There is, however, by now a bulk of evidence suggesting that trustors often do in fact place trust and that trustees regularly deviate from her selfish strategy and honor the trust. To better characterize such empirical findings, one could assume a game of incomplete information with different types of trustees: The trustor knows the set of possible payoff functions, yet she does not have information about the payoff function of the respective trustee she is facing. In this case, she has to base her decision on her subjective prior belief about the proba-

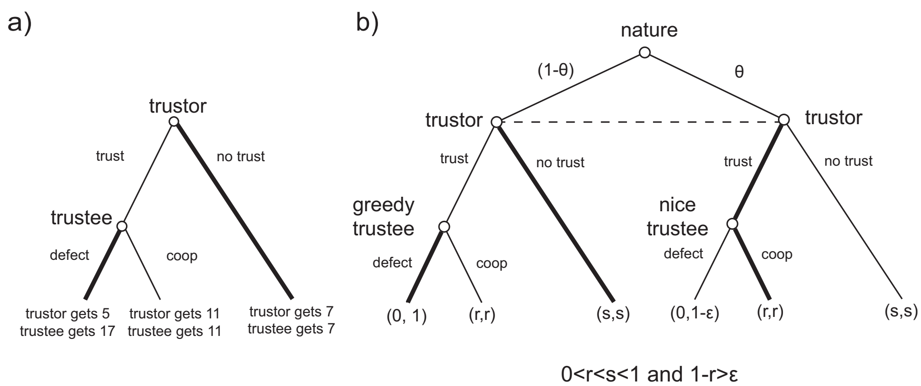

Figure 2 The trust game implemented in this study with payoffs in Euro (a), and the bayesian trust game of incomplete information (b).

bility to meet a trustworthy trustee.[1]

The trust game with two types of trustees and incomplete information can be modeled as a bayesian trust game. There are three types of players: the trustor denoted by λ, and two types of trustees denoted by $\lambda_i, i \in \{g,n\}$, where the index g denotes a greedy trustee (who never returns trust) and n a nice trustee (who always returns trust). In each repetition t of the game, λ is facing a new λ_i, randomly drawn from a population of g,n, where the true composition θ of the population with respect to g,n is unknown to λ. However, λ holds subjective beliefs $p(\theta)$ about the composition of the population.

λ plays finitely many repetitions of the sequential game in the following order: First, nature determines the type of the second mover $\lambda_{g,n}$ with probability θ. Then, λ chooses her action $C_1 = \{trust, no\,trust\}$ based on her belief $p(\theta)$ about the composition the population w.r.t types at repetition t. Next, λ_i chooses his action $C_2 = \{coop, defect\}$ depending on his type, and finally, both players learn their payoffs (see Figure 2b).[2]

For analytical convenience, we assume that the trustor's prior belief about the types in the population follows a beta distribution, which is commonly used to model uncertainty of prior beliefs. The density function of the beta distribution is defined as

$$p(\theta|c_p,d_p) = \frac{\theta^{(c_p-1)}(1-\theta)^{(d_p-1)}}{B(c_p,d_p)} \tag{1}$$

where c_p, d_p are positive shape parameters of the beta distribution, mapping θ to the unit interval and $B(c_p,d_p)$ is the beta function given by

$$\int_0^1 \theta^{(c_p-1)}(1-\theta)^{(d_p-1)}d\theta.$$

[1] Costa-Gomes, Huck and Weizsäcker (2012) provide experimental evidence that initial priors can in fact be considered causal for trust decisions.

[2] We assume that $p(coop|\theta) = \theta$ and $p(defect|\theta) = 1-\theta$.

In the absence of “real” observations, c_p and d_p can be viewed as *pseudocounts*: They are treated *as if* one had observed c_p times cooperation and d_p times defection.

The expected value of the beta distributed $p(\theta)$ with parameters c_p and d_p is given by

$$\overline{\theta} = E(\theta|c_p, d_p) = \frac{c_p}{c_p + d_p}, \tag{2}$$

which is equivalent to the expected probability of observing a trustworthy trustee in the next subgame.

2.1 Equilibrium analysis

Since $r < 1$, the greedy trustee λ_g will choose defect and not return the trust in case player λ chose trust, leaving the former with a payoff of 1 and the latter with a payoff of 0, which is less than λ would have gained if she had chosen not to trust. The nice trustee λ_n, on the other hand, prefers to return the trust over not returning ($1 - \varepsilon < r$), such that λ can gain $r > s$ if she chooses to trust instead of not to trust and meets a player of type λ_n. Thus, λ prefers not to trust over trust if nature draws λ_g, but prefers trust over no trust in the other case. Consequently, player λ's move depends on his belief $p(\theta)$ about the distribution of types in the population. The trustor will trust if his expected utility $u_\lambda(\cdot)$ for this action is greater than the utility for not trusting:

$$\begin{gathered} u_\lambda(\mathit{trust}) > u_\lambda(\mathit{no\,trust}) \\ \Leftrightarrow \\ \overline{\theta} \cdot r + (1 - \overline{\theta}) \cdot 0 > s \\ \Leftrightarrow \\ \overline{\theta} > \frac{s}{r}. \end{gathered} \tag{3}$$

Based on Equation 2 and 3, λ chooses to trust if and only if

$$\frac{c_p}{c_p + d_p} > \frac{s}{r}.$$

Now let us assume λ has observed n choices of randomly chosen λ_i. Of this n choices, the trustor has observed the trustees having c_o times chosen to return trust and $n - c_o$ times chosen not to return the trust. By applying Bayes rule

$$p(\theta|c_o, n) = \frac{p(c_o, n|\theta)p(\theta)}{p(c_o, n)},$$

we can derive the posterior density function for λ as

$$p(\theta|c_o, n) = \frac{\theta^{(c_p + c_o - 1)}(1 - \theta)^{(d_p + (n - c_o) - 1)}}{B(c_p + c_o, d_p + n - c_o)}. \tag{4}$$

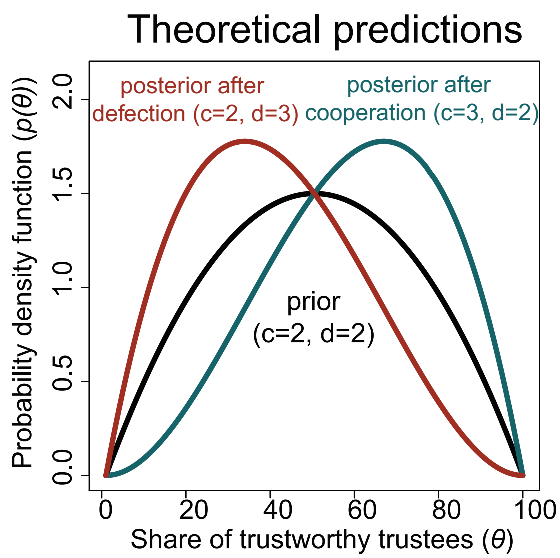

Figure 3 Predicted direction of bayesian updating after observing cooperation (green) and defection (red).

Hence, if the prior distribution of θ is $beta(c_p, d_p)$, then the posterior for θ is $beta(c_p + c_o, d_p + n - c_o)$. As in equation 2, we can now calculate the posterior mean of $p(\theta)$ by inserting the revised shape parameters as

$$\overline{\theta} = \frac{c_p + c_o}{c_p + c_o + d_p + n - c_o} = \frac{c_p + c_o}{c_p + d_p + n}. \tag{5}$$

Equation 5 expresses the trustors' predicted probability that the next trustee will return the trust as well as the expected share of nice trustees in the population. In equilibrium, the trustor will grant trust if equation 5 satisfies the condition in 3, and not grant trust otherwise. This probability increases in the observed share of trustworthy trustees (c_o/n), and decreases in the share of observed greedy trustees. Figure 3 gives an intuition about the updating process after one observation for an arbitrarily chosen prior.

Two hypotheses are derived from the bayesian trust model. First we have an auxiliary hypothesis about the trustors' belief.

Hypothesis 3 (Updating) *The trustor's posterior belief that the next trustee returns trust is higher (lower) if the previous trustee is observed (not) to return trust.*

The main interest we have considering the trustors behavior, however, is whether or not the trustor will factor in friendship homophily. If homophily with respect to trust exists, the choices of two friends should be more likely to be the same as compared to two random strangers. In this case, actions previously observed by friends of the trustee should be considered more informative than the ones by strangers, and should be weighted accordingly. According to the bayesian trust model, the trustor will assign the same probability that the next trustee will return the trust, irrespectively of whether the this trustee is a friend of the previous trustee or not. As an alternative to the null hypothesis of no effect we have the following research hypothesis:

Hypothesis 4 (Anticipated Homophily) *Trustors update their beliefs more strongly when facing two persons from the same friendship network as compared to two persons from different friendship networks.*

3 Experimental Design

The experiment was programmed with the software z-Tree Fischbacher (2007) and conducted in May 2012 in the experimental laboratory of the Max Planck Institute of Economics located in Jena (Germany). A total of 237 undergraduate students (142 female, 95 male) from a wide range of academic disciplines were recruited using the ORSEE software (Greiner, 2004). Each session lasted for approximately one and a half hour and consisted of 27-30 participants with an average payment of 12.72 Euro ($\sigma =$ 4.12, min=7.50, max=19.50), including a show-up fee of 2.50 Euro. The instructions to the experiment were read out aloud and understanding of the rules was ensured by a set of control questions.

In the experiment a trustor plays two consecutive trust games with different trustees (see Figure 2a for the corresponding payoffs in Euro implemented in the experiment). The wording during the experiment was held neutral, e.g. "left" and "right" instead of "trust" and "no trust". Note that the trustor's trustfulness cannot unambiguously be inferred from his choice in the second trust game. More specifically, the decision not to grant trust in the second trust game can be motivated by two factors: It could be a reaction to reduced trustfulness because the trust granted was not returned by the first trustee, or it could be a hostile act of punishment towards the second player for the misbehavior of the first. Therefore we also elicited the trustor's belief about the trustees behavior with the quadratic scoring rule (see Nyarko and Schotter, 2002).[3] The prior beliefs were elicited after the first trust game but of course before informing the trustor about the trustees' behavior. The second trust game followed thereafter. Posterior beliefs were similarly elicited after the second trust decision. See Figure 4 for a timeline of the experiment.

In order to collect data from all trustees irrespective of the trustors' choices, we decided to implement the strategy method (Selten, 1967) before the trustee learns about the trustor's choice.[4]

[3]The trustor is endowed with 11 ECU and is asked to indicate his belief about whether the trustee will return trust or not on a scale from 0 (certainly not return) to 100 (certainly return). Depending on the accuracy of his/her belief, he/she will be paid in an incentive compatible way. To help the subjects to understand the task, the decision screen presented a table with all possible choices and the corresponding payoffs next to each other. The possible pay-offs ranged from 5 ECU to 11 ECU if he wrongly/correctly predicts the trustees' behavior.

[4]In a trust game experiment, it is possible to elicit choices using either the so-called game or strategy method. While the two methods yield similar rates of trustfulness, the strategy method reveals a significantly lower rate of trustworthiness (Casari and Cason, 2009). For a more general discussion see Rauhut and Winter (2010).

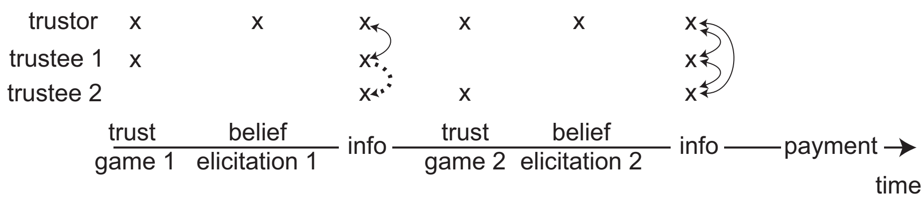

Figure 4 Time line of the experiment. Solid arrows depict the flow of information in all treatments, dashed arrows the flow of information in the information treatments.

The payoff relevant choice for the trustor was randomly chosen from the four decisions (two action decision in the trust games and two for the beliefs). Paying only one decision at random mitigates wealth effects as well as hedging across games. The two trustees were paid according to the outcome of the trust game that they played.

3.1 Experimental Treatments

The experiment facilitated the 2×2 design (see Figure 5). It consists of two consecutive trust games where a group is formed by one trustor, playing with two trustees in sequential order. The trustees in each group were either friends or anonymous to each other. In two of our treatments the trustees did not know each other, while in the other two they were friends. But note that when recruiting subjects we invited 10 subjects to serve as trustors, while we asked another 10 subjects to bring a same-gender friend in all the treatments. [5] The latter 20 subjects where always assigned to the role of the trustees, but depending on the treatment, they were matched to the same or to different trustors. Hence, irrespectively of the treatments the trustees always came to the lab with a friend. Thus, whether or not the trustees' were friends was exogenously manipulated in our experiment. The second manipulation concerned the flow of information. Either the first trustee's choice was known or not known to the second trustee before making his or her choice. Group composition and informational feedback were announced before handing out the instructions for the second trust game in all the treatments. In the information strangers treatment (IS), a group consisted of strangers only. In the information friends treatment (IF), two friends were assigned to the role of the trustees, while those who came without a friend were always assigned to the role of the trustor. The two remaining treatments, the no information strangers (NIS) and the no information friends (NIF) treatment differ from the first two in that no information was transmitted to the second trustee about the first trustee's behavior.

We believe that homophily among friends is an interesting object of study, though we should mention that some caveats to internal validity have to be dealt with. While it is indispensable to invite real-life friends when studying this kind of homophily, it also goes along with a non-negligible loss of control, in particular before and after the experimental sessions. Both trustees may ex ante have agreed to balance their pay-

[5] One could argue that a male bringing a female friend or vice versa would want to impress his or her different-gender counterpart. To eliminate this concern as far as possible, we explicitly asked everyone to bring a same-gender friend.

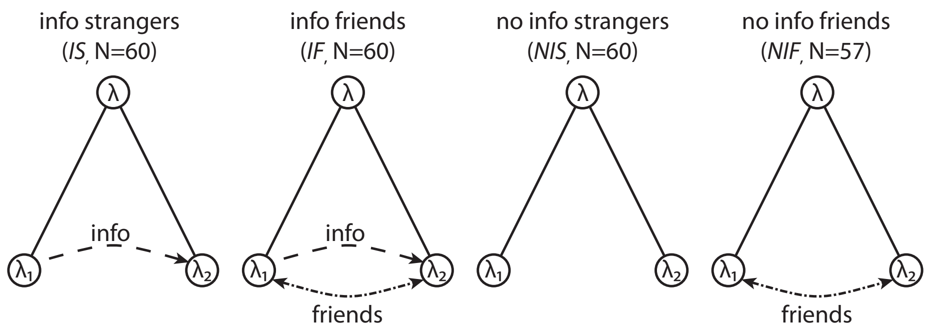

Figure 5 Experimental treatments

ments after the experiment. The first trustee may then cooperate, hoping that his friend gets a chance of unilateral defection. To prevent such behavior, we did not announce the second trust game in detail until the first trust game was completed, such that the first trustee could not foresee the future consequences of his choices. Moreover, the instructions for the first part were the same in all treatments in order to prevent confounding differences in the treatments (see the instructions in the Appendix). Furthermore, friends were seated in distant cubicles and we controlled that no communication was taking place between any of the subjects. Finally, the participants were payed individually and virtually all of them immediately pocketed the money, suggesting that they really saw it as their money, not as the team's income.

Another possible confound is that trustees' image concerns could confound our results if trustees would return the trust in the friends treatments, and withhold it in the stranger treatments, to make a nice impression on their friends. In order to keep the feedback comparable across treatments, trustees in all treatments were informed about their friend's as well as one random stranger's decision at the end of the experiment.

4 Results

We will structure the analysis of our results in three parts: We start by investigating homophily in trustees' choices with and without information, continue with the trustors' beliefs and conclude by analyzing their reaction to (un-)trustworthy behavior.

Result 1 (Conditional Homophily) *Observing a friends' behavior has a stronger effect on a trustee's choice than observing a random stranger's behavior.*

Figure 6 reports the probabilities that two trustees either both return trust or both do not return trust.[6] Learning about a previous trustee's choice results in an about 30 percentage points higher probability of copying the previous trustee only if this first and second trustee are friends. Applying χ^2-tests, we find significant differences between

[6]In the NIS treatment, ten of the subjects came together with a friend and they do not get feedback about their friend's behavior which allows us to include data of these sessions to test if friends behave more similar. This leaves us with 40 friend-couples in the no information condition.

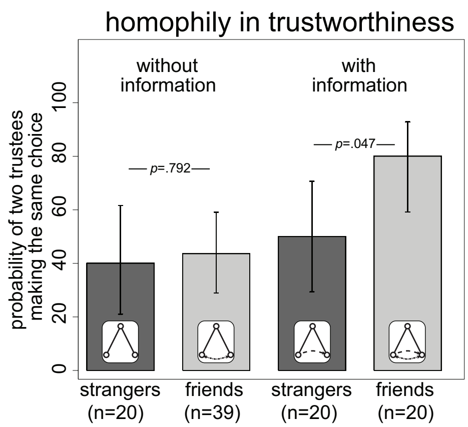

Figure 6 Homophily in trustworthiness among friends. Probability that two trustees make the same choice, depending on whether they are friends or not, and whether the second trustee learns about the first trustee's choice or not. The *p*-values result from treatment-to-treatment χ^2-tests. All other comparisons are not significantly different for any conventional level (see also Table 2 in the Appendix).

the IF-treatment and the IS-treatment (χ^2-test, $p = .047$, $n = 40$). The finding that respondents are only affected by observing their friends' choices gives support to our "Conditional Trust Hypothesis", which conjectured a stronger effect of information on friends as compared to strangers.

Result 2 (Unconditional Homophily) *There is not a higher correlation in choices among trustees that are friends than trustees that are strangers if the trustees in a group cannot observe each other's behavior.*

In the no information treatments, the probability to make the same choice is not significantly different between friends and strangers (χ^2-test, $p = .792$, $n = 59$). We thus do not find support for our "Unconditional Homophily Hypothesis", which stated that trustworthiness is shared among friends *unconditionally* of the knowledge of the friend's behavior.

Result 3 (Updating) *Trustors update their posterior beliefs consistent with their observation: Observing trustworthiness increases their posterior beliefs of trustworthiness in the population, observing defection decreases posterior beliefs.*

The right panel in Figure 7 plots the observed updating behavior in the experiment. The prior of cooperation (black line) plots the estimated distribution of the first belief elicitation, elicited before learning about the first trustee's choice. The posterior (red and green line, respectively) plots the estimated distribution *after* learning about the first trustee's choice, but before learning about the second's choice.[7] We observe that

[7]Due to the rather small sample size, we pooled the observed updating data across treatments for this figure.

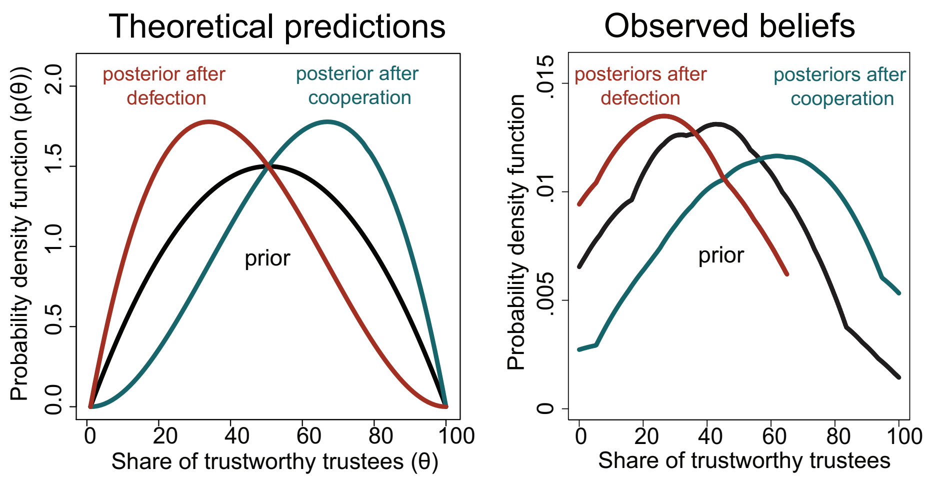

Figure 7 Bayesian updating. Theoretical predictions of bayesian updating after trust is (not) returned (left) and separate kernel density estimates for priors and posteriors after trust is (not) returned (right).

updating is consistent with the "Updating Hypothesis" derived in Section 2: Observing trust being returned increases the posterior, observing withheld trust decreases it.

This observation is supported in an OLS-regression analysis (see Table 1).[8] Here we find a strong and significant effect of feedback: Controlling for initial priors, we find that trustors who observed returned instead of withheld trust in the first trust game have a 27 percent higher posterior that the next trustee will cooperate as well.

Result 4 (Anticipated Homophily) *Observing two trustees' that are friends has the same effect on the trustors' posterior beliefs as observing two trustees that are not friends.*

Our data does not give support for Hypothesis 4. Rather, the right model reported in Table 1 confirms that the posterior beliefs depend on the prior beliefs and are updated based on the trustees' behavior only. The predicted treatment differences after observing cooperation or defection are insignificant and do not add additional explanatory power.

Result 5 (Indirect Tit for Tat) *Uncooperative behavior of a a trustee often leads to retaliation towards this trustee's friend.*

Consistent with the trustors' beliefs we observe that the trustor trusts the second trustee less if the first trustee did not return the trust. However, since the trustors' posterior beliefs do not differ whether or not the trustees' are friends (see result 4), it follows from our model that their decision to trust should not differ between friends and strangers either. However, we observe that the trustors' do not behave in line with

[8] A tobit estimation give qualitatively similar results.

Table 1 Determinants of Posterior Beliefs

	posterior (cooperation)	posterior (cooperation)
prior (cooperation)	0.558***	0.556***
	[0.385,0.731]	[0.375,0.737]
feedback (cooperation = 1)	27.74***	21.71*
	[19.17,36.30]	[4.563,38.87]
info stranger (IS)		ref.
info friends (IF)		-5.252
		[-21.19,10.69]
no info strangers (NIS)		-8.418
		[-24.05,7.215]
no info friends (NIF)		-7.285
		[-23.24,8.673]
cooperation × IF		4.883
		[-19.73,29.49]
cooperation × NIS		6.305
		[-18.49,31.10]
cooperation × NIF		12.55
		[-12.43,37.52]
constant	5.762	11.25
	[-2.608,14.13]	[-2.050,24.56]
N	79	79
R^2	0.542	0.554
BIC	696.18	720.28

95% confidence intervals in brackets

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

OLS-Regressions estimating the determinants of posterior beliefs, i.e. the trustor's belief about the second trustee's choice after the trustor was informed about the first trustee's choice. prior(cooperation) measured the belief that the first trustee will return trust before learning about his/her choice, feedback is a dummy taking the value 1 if the first trustee returned the trust, IF,NIS, and NIF are dummies taking the value 1 in the Information Friends, No Information Strangers, and No Information Friends treatment, respectively, and cooperation × IF(NIS,NIF) are the respective interaction effects.

their beliefs and instead are more likely to trust if the trustee's friend has shown to be trustworthy in the first trust game, compared to if the trustees' are strangers to each other (see Figure 8 a)).[9]

Likewise, they are less likely to trust if the trustees' friend has shown to be untrustworthy compared to if the trustees were strangers. Figure 8 b) shows that the trustor is significantly more likely to mirror the first trustees decision by (not) placing trust in the second trustee if the first (did not) proved to be trustworthy when the trustees' are friends instead of strangers. The trustors answer cooperative behavior with coopera-

[9]We refer to Table 3 in the appendix for the statistical comparisons.

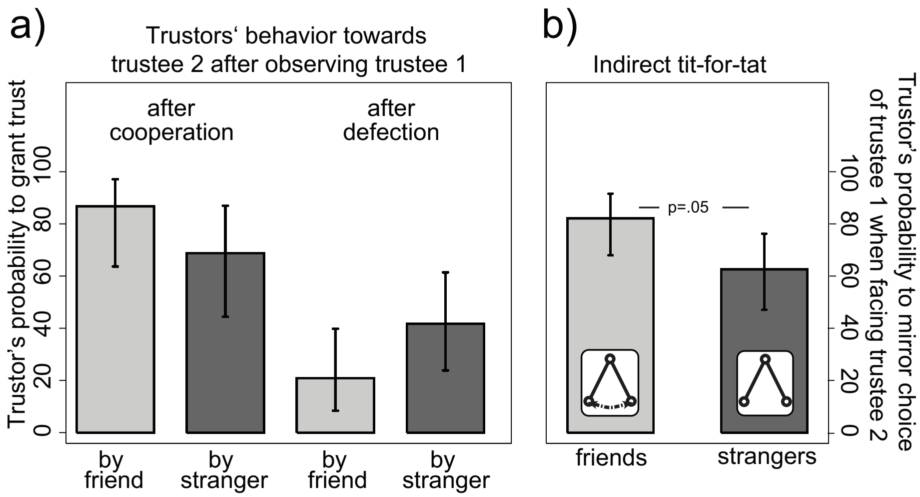

Figure 8 Behavioral reactions to (un-)trustworthy behavior. Trustors' average trust rates towards the second trustee after cooperation and defection by the first trustee if the second is a stranger or a friend of the first (left), and the share of trustors playing indirect tit-for-tat by mirroring the first trustee's choice to cooperate or defect towards the second trustee (right). The *p*-value in the right figure is obtained by a χ^2-test, additional comparisons for the left panel can be found in Table 3 in the Appendix.

tion and non-cooperative behavior with defection *indirectly*, targeting it towards some member of the first person's friendship network instead of the initial player. This behavior may accordingly be termed *Indirect Tit-For-Tat* in the sense of Axelrod (1984).

5 Conclusion

In this paper we study the phenomenon of homophily with respect to trust, defined as the tendency for people to have ties with people who are similarly trustworthy as themselves. Even though lab-experimental research with friends may require a carefully considered design to ensure internal validity, we believe our questions to be important and interesting. We ask whether trustworthiness towards strangers is shared among friends in a network, and our results are affirmative under specific conditions: If people can observe their friend's behavior towards outsiders prior to making their own decisions, they are likely to imitate this behavior. Imitation occurs significantly less if people observe the behavior of other participants not known to them. Our results support emergence of homophily with respect to trust. Since contacts to members of the same network are not only more influential, but usually also more frequent than contacts to outsiders, people should become more similar within networks than across networks.

On the other hand, we asked whether trustors anticipate homophily, and thus condition on their knowledge about the friends of potential exchange partners. Here, the answer is no: When asking the trustors for their beliefs whether the next trustee will be trustworthy or not, it made no difference whether the trustees where friends or strangers. When observing the trustors' actions rather than their beliefs, however, they very well condition on previous encounters from the same network. Untrustworthy behavior by one person is much more likely to be reciprocated through a trustor if both trustees are friends, just as trustworthy behavior is more likely to be honored if the trusted parties come from the same friendship network. This behavior may be called Indirect Tit-For-Tat: Trustors reciprocate the previous behavior (as in Tit-For-Tat (Axelrod, 1984)), but it is an indirect form of punishments and rewards, since it is not targeted towards the defector or cooperator himself, but to someone from his network. Our results suggest that agents should be careful to preserve not only their own reputation but the reputation of the network they choose to be part of. Future research should look at this phenomenon more carefully.

6 Acknowledgements

We thank Vittoria Maria Levati and Michael Trost for valuable comments to the manuscript, and Ria Stangneth, Lisa Schöttl, Dvin Galstian Pour and Andreas Matzke for valuable research assistance. No IRB approval is required at the authors' institution. The usual disclaimer applies.

References

Arrow, K.J. 1974. *The limits of organization.* Vol. 1, Norton.

Axelrod, Robert. 1984. *The evolution of cooperation.* New York:Basic Books.

Baccara, Mariagiovanna, and Leeat Yariv. 2013. "Homophily in Peer Groups." *American Economic Journal: Microeconomics*, 5(3): 69–96.

Bott, H. 1928. "Observation of play activities in a nursery school." *Genetic Psychology Monographs*, 4(1): 44–88.

Casari, Marco, and Timothy N. Cason. 2009. "The strategy method lowers measured trustworthy behavior." *Economics Letters*, 103(3): 157 – 159.

Coleman, James Samuel. 1990. *Foundations of Social Theory.* Cambridge, Mass.:Belknap Pr.

Costa-Gomes, Miguel A., Steffen Huck, and Georg Weizsäcker. 2012. "Beliefs and actions in the trust game: Creating instrumental variables to estimate the causal effect." WZB Discussion Paper SP II 2012-302, Berlin:WZB.

Dasgupta, Partha. 1988. "Trust as a Commodity." In *Trust.* , ed. Diego Gambetta, Chapter 4, 49–72. Blackwell Publishers.

DiPrete, Thomas A., Andrew Gelman, Tyler McCormick, Julien Teitler, and Tian Zheng. 2011. "Segregation in Social Networks Based on Acquaintanceship and Trust." *American Journal of Sociology*, 116(4): pp. 1234–83.

Fischbacher, Urs. 2007. "Zurich Toolbox for ready-made economic experiments." *Experimental Economics*, 10(2): 171–178.

Goeree, Jacob K., Margaret A. McConnell, Tiffany Mitchell, Tracey Tromp, and Leeat Yariv. 2010. "The 1/d Law of Giving." *American Economic Journal: Microeconomics*, 2(1): pp. 183–203.

Golbeck, Jennifer. 2009. "Trust and nuanced profile similarity in online social networks." *ACM Trans. Web*, 3(4): 12:1–12:33.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology*, 78(6): pp. 1360–1380.

Greiner, Ben. 2004. "An Online Recruitment System for Economic Experiments." In *GWDG Bericht 63*. , ed. Kurt Kremer and Volker Macho, 79–93. Göttingen:Ges. für Wiss. Datenverarbeitung.

Hatfield, Elaine. 1984. "The dangers of intimacy." In *Communication, intimacy, and close relationships*. , ed. Valerian J. Derlega, 207–220. Academic Press.

Igarashi, Tasuku, Yoshihisa Kashima, Emiko S. Kashima, Tomas Farsides, Uichol Kim, Fritz Strack, Lioba Werth, and Masaki Yuki. 2008. "Culture, trust, and social networks." *Asian Journal of Social Psychology*, 11(1): 88–101.

Jackson, Matthew O, and Alison Watts. 2002. "The evolution of social and economic networks." *Journal of Economic Theory*, 106(2): 265–295.

Jackson, Matthew O., and Brian W. Rogers. 2007. "Meeting Strangers and Friends of Friends: How Random Are Social Networks?" *The American Economic Review*, 97(3): pp. 890–915.

Kataria, Mitesh, and Fabian Winter. 2013. "Third party assessments in trust problems with conflict of interest: An experiment on the effects of promises." *Economics Letters*, 120(1): 53 – 56.

Lazarsfeld, P. F., and R. K. Merton. 1954. "Friendship as a social process: A substantive and methodological analysis." In *Freedom and Control in Modern Society*. , ed. M. Berger, T. Abel and C. Page, 18–66. New York:Van Nostrand.

Leider, Stephen, Markus M. Möbius, Tanya Rosenblat, and Quoc-Anh Do. 2009. "Directed Altruism and Enforced Reciprocity in Social Networks." *The Quarterly Journal of Economics*, 124(4): 1815–1851.

Leider, Stephen, Tanya Rosenblat, Markus M. Möbius, and Quoc-Anh Do. 2010. "What do we Expect from our Friends?" *Journal of the European Economic Association*, 8(1): 120–138.

Loomis, Charles P. 1946. "Political and Occupational Cleavages in a Hanoverian Village, Germany: A Sociometric Study." *Sociometry*, 9(4): pp. 316–333.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology*, 27(1): 415–444.

Nyarko, Yaw, and Andrew Schotter. 2002. "An Experimental Study of Belief Learning Using Elicited Beliefs." *Econometrica*, 70(3): 971–1005.

Rauhut, Heiko, and Fabian Winter. 2010. "A Sociological Perspective on Measuring Social Norms by Means of Strategy Method Experiments." *Social Science Research*, 39(6): 1181 – 1194.

Selten, Reinhard. 1967. "Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperimentes." In *Beiträge zur experimentellen Wirtschaftsforschung*. , ed. Heinz Sauermann. Tübingen:J. C. B. Mohr (Paul Siebeck).

Tang, Jiliang, Huiji Gao, Xia Hu, and Huan Liu. 2013. "Exploiting homophily effect for trust prediction." *WSDM '13*, 53–62. New York, NY, USA:ACM.

Appendix

A Additional statistical tests

	NIS (8/20)	NIF (17/39)	IS (10/20)
NIF (17/39)	.792	–	–
IS (10/20)	.525	.64	–
IF (16/20)	.01	.008	.047

Table 2 *P*-values for treatments-wise comparison of the probability that two trustees make the same choice obtained from χ^2-tests (see also Figure 6). The numbers in parenthesis display the number of cases where two trustees made the same choice and the total number of observations for each case, respectively.

	friend cooperated (13/15)	stranger cooperated (11/16)	friend defected (5/24)
stranger cooperated (11/16)	.233	–	–
friend defected (5/24)	$< .001$	.002	–
stranger defected (10/24)	.005	.093	.119

Table 3 *P*-values for treatments-wise comparison of the probability to grant trust after cooperation/defection by the first trustee obtained from χ^2-tests (see also Figure 8). The numbers in parenthesis display the number of cases where trust was granted after cooperation/defection of a friend/stranger to the first trustee and the total number of cooperation/defection by the first trustee, respectively.

B Instructions (English translation)

Instructions
(equivalent for all treatments)

Welcome and thank you for participating in this experiment! Please read these instructions carefully. They are identical for all participants. For being in time you will receive a show-up fee of 2.50 Euro. During the experiment you will have the opportunity to earn additional money. How much depends on your decisions in the experiment, as well as on the decisions of other participants. All amounts in the experiment will be displayed in ECU (experimental currency units). All earned ECU will be exchanged into Euro at the end of the experiment and disbursed in cash according to the following exchange rate:

1 ECU = 1 Euro

Please avoid talking to your neighbors from now on. Switch off your mobile and remove all unneeded things from your desk. It is important that you act according to these rules, as otherwise we will have to exclude you from the experiment and therefore from any payment as well. In case of questions please raise your hand and we will answer them individually.

For the following experimental session we invited a total of 30 participants. We asked some of you to bring a male friend. **However, note that friends will never directly interact with each other**.

The experiment consists of two parts. Each part consists of two different tasks. We will first explain the rules for the first part, and after the part has been finished by all participants, we will explain the rules for the second part. Some of you will participate only in part one, others only in part two and some of you will participate in both. The decisions made in part 1 cannot directly affect earnings in part 2 and vice versa. You will be told in which part you are selected to participate in due time. In the first part of the experiment you will be matched in groups of two. Remember that the person you will be matched with will not be your friend.

Part 1A
Decision Task

The game involves two positions, namely the position of **Player 1** and the position of **Player 2**. The game starts with Player l's move "left" or "right". If Player 1 chooses "right", the game is over. Only after Player l's move "left", Player 2 has to choose between his move "left" or "right", which ends the game.

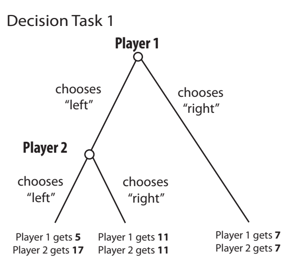

The decision process is graphically illustrated above. At the top decision node, the origin of the game tree, Player 1 must decide whether to use "left" or "right". Only after the move "left", Player 2's choice between "left" and "right" matters. The payoffs are given at the endpoint with the upper component being Player l's monetary gains, whereas the lower component is Player 2's payoff. Thus if Player 1 chooses "right", both Players earn 7 ECU. In case of Player l's choice "left", the payoffs depend on Player 2's move: If Player 2 uses "left", Player 1 receives 5 ECU and Player 2 receives 17 ECU, whereas they both earn 11 in case of Player 2's moves "right".

When playing the game, we will ask Player 2 how he would decide in case that Player 1 chose "left". Thus Player 2 must choose between "left" and "right" without being informed how Player 1 has decided. The reason for this is that we want to know how Player 2 would have reacted to "left" even if Player 1 had chosen "right". So when deciding between "left" and "right" as Player 2, he should imagine that Player 1 has chosen "left"

Part 1B
Prediction Task

Prediction Task

The prediction task will only be performed by Player 1. Before Player 1 learns about Player 2's decision, we will ask him how likely he thinks it is that Player 2 played option "left" (and "right" respectively). The better Player 1 predicts Player 2's choice, the more Player 1 can earn. No one beside Player 1 can earn money in this task, and Player 1's decisions do not affect the other person's payoffs.

Predicting Player 2's Choice

After Player 1 made his choice of option "right" or "right", he will be given an opportunity to earn money by predicting Player 2's choices of "left" or "right" in game 1. He will be asked to predict Player 2's choice by stating how likely he thinks Player 2 will choose either one of his two choices – "left" or "right".

To make a prediction we will supply you with a prediction table as follows:

Prediction Table:

Probability that Player 2 will choose "left"	Probability that Player 2 will choose "right"	Your payoff in case Player 2 chose "left"	Your payoff in case Player 2 chose "right"
20	80	7.16	10.76
19	81	7.06	10.78
18	82	6.97	10.18
17	83	9.87	10.83
16	84	6.77	10.85
15	85	6.67	10.86
14	86	6.56	10.88
13	87	6.46	10.90
12	88	6.35	10.91
11	89	6.25	10.93
10	90	6.14	10.94
9	91	6.03	10.95
8	92	5.92	10.96
7	93	9.81	10.97
6	94	5.70	10.98
5	95	5.58	10.98

OK

This table allows Player 1 to make a prediction **of Player 2's choice** by highlighting a line in the table indicating what the chances are that Player 2 will choose "left" or "right". For example, say Player 1 thinks there is a 10% chance that Player 2 will choose "left", and hence a 90% chance that "right" will be chosen. This indicates that Player 1's believes that "left" is less likely to be chosen than "right" by a considerable margin. If this is Player 1's belief about the likely choice of Player 2, then he simply highlights the line containing the 10 in the left column and confirms by clicking the OK-button. After that, the computer will look at the choice actually made by Player 2 and compares his choice to Player 1's predictions. Player 1 will then be paid for his prediction as follows:

Suppose Player 1 predicts that Player 2 will choose "right" with a 90% chance and "left" with a 10% chance. In that case he will highlight the line indicating a 10% for option "left" and submits. Suppose now that Player 2 actually chooses "left". In that case Player 1's earnings will be

Prediction Earnings = 11 ECU- 6 ECU * [$(1\text{-}.01)^2$- $(0.90)^2$] = 6.14 ECU.

The payoff is calculated by giving Player 1 a fixed amount of 11 ECU from which we will subtract an amount which depends on how inaccurate his prediction was. To do this when we find out what choice Player 2 has made we will we will take the number Player 1 assigned to that choice, in this case 10% on "left", subtract it from 100% and square it. We will then take the number Player 1 assigned to the choice not made by Player 2, in this case the 90% assigned to "right", and square it also. These two squared numbers will then be multiplied by 8 and subtracted from the 11 ECU we initially gave Player 1 to determine his final earnings.

Note that the worst Player 1 can do under this payoff scheme is to state that he believes that there is a 100% chance that a certain action is going to be taken and assign 100% to that choice when in fact the other choice is made. Here his payoff from prediction would be 5 ECU. Similarly, the best he can do is to guess correctly and assign 100% to that choice which turns out to be the actual choice chosen. Here Player 1's payoff will be 11 ECU.

Most importantly to remember is that the best thing Player 1 can do to maximize the expected size of the prediction earnings is to simply state his true beliefs about what he thinks Player 2 will do. Any other prediction will decrease the amount he can expect to earn as a prediction earning.

Who will learn about your choices in Decision Task 1?

Player 1's decision in part 1 will be revealed to Player 2 and Player 2's decision will be revealed to Player 1. At some point in the experiment, those participants who came to the lab with a friend will be told about the role and decision of their respective friend. In addition to that, another person will be informed about Player 2's decision and Player 2 will be informed about this random person's decision.

Your earnings in part 1

Player 1 will make some additional decisions in the next part. One of all his choices will randomly selected to determine his payment. Player 2's earning is determined by his choice and Player 1's choice in the decision task in part 1.

Player 3 will not earn anything in part 1

Do you have any questions so far? If you have no more questions, please click OK on the computer screen.

Part 2A
(differences between treatments in brackets)
Decision Task

This is Part 2 of the experiment. Part 2 is similar to part 1.

The person who was called Player 1 in part 1 will again be in the role of Player 1. Player 1 will now interact with **another** person, with whom he had not interacted before. This person will be called **Player 3**. Player 3 has not made any decision in part 1, and Player 2 from the previous part will not make any decisions in part 2. [**NIS:** Player 3 is neither a friend of Player 1 nor of Player 2.] [**IS:** Player 3 is neither a friend of Player 1 nor of Player 2.] [**NIF:** Player 3 is a friend of Player 2.] [**IF:** Player 3 is a friend of Player 2.]

The decision task involves two positions, namely the position of **Player 1** and the position of **Player 3**. The game starts with Player l's move "right" or "left". If Player 1 chooses "left", the game is over. Only after Player l's move "right", Player 3 has to choose between his move "left" or "right", which ends the game.

The decision process is graphically illustrated above. At the top decision node o, the origin of the game tree, Player 1 must decide whether to use "right" or "left". Only after the move "right", Player 3's choice between "left" and "right" matters. The payoffs are given at the endpoint with the upper component being Player l's monetary win, whereas the lower component is Player 3's payoff. Thus if Player 1 chooses "left", both Players earn 7 ECU. In case of Player l's choice is "right", the payoffs depend on Player 3's move: If Player 3 uses "left", Player 1 receives 5 ECU and Player 3 receives 17 ECU, whereas they both earn 11 in case of Player 3's move "right".

When playing the game, we will ask Player 3 how he would decide in case that Player 1 chose "left". Thus Player 3 must choose between "left" and "right" without being informed how Player 1 has decided. The reason for this is that we want to know how Player 3 would have reacted to "left" even if Player 1 has chosen "right". But, of course, Player 3's decision matters only for the monetary payoffs if Player 1 chooses "left". So when deciding between "left" and "right" as Player 3, you should imagine that Player 1 has chosen "left".

IS: Before Player 3 makes his decision, he will be informed about Player 2's choice in Part 1A.

IF: Before Player 3 makes his decision, he will be informed about Player 2's choice in Part 1A.

Part 2B

Prediction Task

The prediction task will again only be performed by Player 1. Before Player 1 learns about **Player 3's decision**, we will ask him how likely he thinks it is that **Player 3** played option "left" (and "right" respectively). The procedure follows the exact same rules as explained for the prediction task in part 1.

Remember that it is in your best monetary interest to state your beliefs truthfully.

Who will learn about your choices in Decision Task 2?

Player 1's decision in part 1 will be revealed to Player 3 and Player 3's decision will be revealed to Player 1. At the very end of the experiment, those participants who came to the lab with a friend will be told about the role and the decision of their respective friend. In addition to that, [**NIF:** another person] [**IF:** another person] [**NIS:** player 2] [**IS:** player 2] will be informed about Player 3's decision and Player 3 will be informed about this [**IS:** other person's] [**NIS:** player 2's] decision.

Your final earnings

Player 3's earning is determined by his choice and Player 1's choice in the decision task in part 2. Player 1 made four decisions (two in part 1 and two in part 2). The computer will randomly select one of these choices to determine his payoff.