Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Trust, but verify? When trustworthiness is observable only through (costly) monitoring

Timo Goeschl
Johannes Jarke

Trust, but verify? When trustworthiness is observable only through (costly) monitoring

Timo Goeschl, Universität Heidelberg

Johannes Jarke, Universität Hamburg

Die Working Paper Series bieten Forscherinnen und Forschern, die an Projekten in Federführung oder mit der Beteiligung der Fakultät Wirtschafts- und Sozialwissenschaften der Universität Hamburg tätig sind, die Möglichkeit zur digitalen Publikation ihrer Forschungsergebnisse. Die Reihe erscheint in unregelmäßiger Reihenfolge.

Jede Nummer erscheint in digitaler Version unter
https://www.wiso.uni-hamburg.de/de/forschung/working-paper-series/

Kontakt:

WiSo-Forschungslabor
Von-Melle-Park 5
20146 Hamburg
E-Mail: experiments@wiso.uni-hamburg.de
Web: http://www.wiso.uni-hamburg.de/forschung/forschungslabor/home/

Trust, but verify? When trustworthiness is observable only through (costly) monitoring

Timo Goeschl[*]   Johannes Jarke[†]

Abstract

For theoretical and empirical reasons, trust is expected to be lower in economic interactions in which trustors can observe trustworthiness only through (costly) monitoring. We examine this conjecture by investigating the impact of a (costly) monitoring environment on trust using data from 152 subjects participating in a modified finite-horizon binary trust game. The three treatment conditions vary observability and the cost of monitoring. We find that compared to perfect observability of trustworthiness, trustors do not trust less when trustworthiness can only be observed through costless or costly deliberate monitoring. When monitoring is costly, the same level of trust is supported by a significantly reduced amount of information on trustworthiness, acquired by trustors mainly in early stages of the repeat interaction. As a result, the efficiency of interactions is not lower when trustworthiness is costly to observe, though the distribution shifts in favor of trustees. (*JEL* C92, C72, D03, D80)

1 Introduction

Trust has been acknowledged as an important prerequisite for realizing the gains from cooperating in the many circumstances in which contractability is limited (Arrow, 1972; Greif, 1993; Zak & Knack, 2001; Bohnet et al., 2001). In such circumstances, contracts will be incomplete and when engaging in transactions that expose them to potential exploitation, parties need to have confidence that the other party will not behave opportunistically. Given its significant role in supporting economic and other social transactions,[1] understanding the presence, nature, and scale of trust as well

---

[*]Alfred-Weber-Institute of Economics, Universität Heidelberg. Bergheimer Strasse 20, D-69115 Heidelberg, Germany. Phone: +49 6221 54 8010. E-mail: goeschl@eco.uni-heidelberg.de.

[†]Corresponding author. School of Business, Economics and Social Sciences, Department of Socioeconomics, Universität Hamburg. Welckerstrasse 8, D-20354 Hamburg, Germany. Phone: +49 40 42838 8768. E-mail: johannes.jarke@wiso.uni-hamburg.de.

[1]For seminal contributions on the role of trust in economic performance see Putnam (1993b), Fukuyama (1995) and Knack & Keefer (1997). For recent literature see Guiso et al. (2004), Annen (2013), Yu et al. (2015), Georgarakos & Fürth (2015), and the references therein. On the relation of trust and governmental performance or community governance see Putnam (1993a), Broix & Posner (1998), DiPasquale & Glaeser (1999), Jackman & Miller (1998), Knack (2002), Bowles & Gintis (2002), Carpenter et al. (2004), Letki & Evans (2005), Berggren & Jordahl (2006), and Bjørnskov (2011).

as its determinants has attracted considerable attention over the years.[2] Considerable progress has come from studies of trust under controlled conditions, such as those based on the «trust game» (or its variants the «investment game» and the «gift-exchange game») in experimental economics (Camerer & Weigelt, 1988; Fehr et al., 1993; Berg et al., 1995).

The (game) theoretical literature highlighted early on that among the determinants of trust, both repetition and informational richness stand out as important structural features of the environment in which incompletely contractible interactions take place (Trivers, 1971; Rubinstein, 1979; Kreps et al., 1982; Fudenberg & Maskin, 1986; Kandori, 1992a,b; Fudenberg et al., 1994). If agents are sufficiently patient, the «shadow of the future» implicit in repeat interactions is predicted to support trust through a disciplining effect on opportunistic behavior. Likewise, informational richness is predicted to support trust by ensuring that non-opportunistic behavior is observable. This assures that agents are able to condition their actions on information about the past actions of their counterparts, hence justifying investment in a reputation for good behavior. These predictions have not only been borne out by laboratory experiments with immediate and complete feedback (Camerer & Weigelt, 1988; Anderhub et al., 2002; Engle-Warnick & Slonim, 2004, 2006a,b; Cochard et al., 2004).[3] More recently, researchers have in fact demonstrated (see section 2 for a review) that the maintenance of trust imposes minimal requirements on the structure of repeat interactions as long as the information environment is sufficiently rich.[4]

The starting point of this paper is the observation that many economic interactions take place in information environments that are not naturally rich. One familiar deficiency is that the trustee's response to being trusted is not automatically and freely observable by the trustor, even when trustor and trustee interact repeatedly in stable pairs. Typical examples are situations of spatially extended trading networks that put physical distance between the trustor and the trustee (Greif, 1993) or situations of structural information asymmetry such as expert knowledge in the health or car repair markets (Wolinsky, 1995; Emons, 1997; Dulleck & Kerschbamer, 2006). In these circumstances, trustors can choose to either stay uninformed or *monitor* the trustee's action. By «monitoring» we mean deliberate action to remedy a lack of information

---

[2]The literature is vast by now. See Ostrom & Walker (2003), Guiso et al. (2008), Sapienza et al. (2007), Fehr (2009a), Johnson & Mislin (2011), and Thöni et al. (2012) as useful entrances.

[3]The general lesson that repetition supports cooperation has also been found with a variety of other experimental social dilemma games, such as the prisoners' dilemma game (e.g. Andreoni & Miller, 1993; Cooper et al., 1996; Dal Bó, 2005; Duffy & Ochs, 2009; Dal Bó & Fréchette, 2011), the public good game (see Andreoni & Croson, 2008, for an overview), and the gift-exchange game (Kirchler et al., 1996; Fehr et al., 1998; Falk et al., 1999; Gächter & Falk, 2002).

[4]This result is mirrored by similar findings in a PD context. Camera & Casari (2009) and Camera et al. (2012) show how cooperation can be maintained in an indefinitely repeated prisoners' dilemma following a strict stranger protocol as long as identities and histories of co-players are public information.

on co-players' actions. The need[5] for as well as the feasibility[6] of monitoring is a routine feature of economic and other social interactions and is particularly salient when monitoring is not casual, but requires costly effort (Ostrom, 1990; Weissing & Ostrom, 1991; Ostrom & Gardner, 1993; Rustagi et al., 2010). The basic logic is also enshrined in the Russian proverb «Доверяй, но проверяй» («trust, but verify»). The question at the heart of this paper is how such limited observability of trustworthiness and the resulting need for (costly) monitoring impacts on trust in repeat interactions.

Is there less trust in interactions in which trustors can observe their co-player's trustworthiness only through (costly) monitoring? Theoretical and behavioral considerations give reason to expect lower trust when repeat interactions have to take place in a (costly) monitoring environment (see section 3.3 for details). Similarly, existing experimental evidence on the impact of exogenously imposed information imperfections in cognate game forms (Sell & Wilson, 1991; Holcomb & Nelson, 1997; Cason & Khan, 1999; Grechenig et al., 2010; Ambrus & Greiner, 2012) points to lower trust. However, there is to our knowledge as yet no experimental evidence that directly addresses this question. Such evidence would not only advance our understanding of the impact of information acquisition costs on trust, but also our understanding of whether and how trustors choose to learn to trust their trustee.[7]

We investigate the impact of a (costly) monitoring environment on trust using data from 152 subjects participating in a modified finite-horizon binary trust game that we term—for short—the *monitoring game*. In the original trust game, a first mover (the «trustor») chooses between an outside option and a trust move that renders her vulnerable to exploitation by a second mover (the «trustee»). The trustee may either reward the trustor's trusting move at a personal cost or exploit the opportunity. When repeated, the trustor knows at the outset of the next round whether her trust was rewarded or not and can decide accordingly. The essential design variation in the *monitoring game* is that trusting first movers do not automatically learn the outcome at the end of a round, i.e. neither the trustee's action nor the payoffs from which that

---

[5]Narratively, efforts to overcome imperfect information on co-players' actions have been recognized in a variety of relevant contexts, such as shared resource management (Ostrom, 1990; Ostrom & Gardner, 1993; Rustagi et al., 2010), production teams (Alchian & Demsetz, 1972; Kandel & Lazear, 1992; Dong & Dow, 1993), labor relations (Shapiro & Stiglitz, 1984; Kanemoto & MacLeod, 1991; Lazear, 1993), finance (Williamson, 1986, 1987; Armendáriz & Morduch, 2005) or neighborhood watch (Sampson et al., 1997). See Ben-Porath & Kahneman (2003), Miyagawa et al. (2008), and Awaya (2014) for theoretical motivations.

[6]An early example is Varian (1990, p. 153) who commented that the agency literature «typically assumes that principals are unable to observe the characteristics or the actions of the agents ... However, in reality, it is often not the case that agents' characteristics or effort levels are really unobservable; rather, they simply may be very costly to observe. One may choose to model high-costs actions as being infeasible actions, but in doing so, one may miss some interesting phenomena.»

[7]This, in turn, has the potential to help identify plausible drivers of agent's trust behavior in the trust game. Whether behavior in the trust game captures trust in an adequate fashion (or just risk seeking or altruism) is a question of concern both to behavioral economists (e.g. Glaeser et al., 2000; Cox, 2004; Eckel & Wilson, 2004; Karlan, 2005; Schechter, 2006, 2007; Ashraf et al., 2006; Ben-Ner & Halldorsson, 2010; Houser et al., 2010; McEvily et al., 2012) and social scientists more generally (e.g. Rousseau et al., 1998; Elster, 2007).

action could be deduced. Instead, the trustor chooses whether to monitor the trustee's action in that round. If she makes an active monitoring decision, he observes the trustee's action. In other words, the trustor can choose to «trust, but verify», but also to trust without such verification.[8] In the latter case, the trustee's action in that round remains hidden «forever» (i.e. until the end of the supergame). While the trustee is aware that the trustor has the monitoring option, he does not learn whether he is actually being monitored.[9] We implement the stage game in a twelve-round repetition.

To study the impact of limited observability and a (costly) monitoring option on trust, we used three treatments. The «Baseline» treatment implemented a standard finite horizon trust game with *perfect observability*, that is, trustors were automatically informed about the trustee's actions without incurring a cost. This replication of previous research (e.g. Anderhub et al., 2002; Engle-Warnick & Slonim, 2004, 2006b,a; Slonim & Guillen, 2010) returned the typical pattern of frequent trust (about two in three cases) until close to the terminal period, and a sharp decline in the final two periods. In the «Costless Monitoring» treatment, trustors were not automatically informed about the trustee's action, but could remedy this limited observability at no cost by taking an active decision to monitor the trustee. While from a theoretical point of view, costless monitoring is indistinguishable from perfect observability, the active decision to monitor involves—from a behavioral perspective—an element of distrust towards the trustee (McEvily et al., 2012).[10] Like the «Costless Monitoring» condition, the «Costly Monitoring» treatment required a deliberate monitoring decision from the trustor in order to observe the trustee's action, but now monitoring involved a cost equal to one third of the per-period gain from rewarded trust.

Comparing behavior in the three treatment conditions, we arrive at three key results. The first is that limited observability does not adversely affect trust if a (costly) monitoring option is present. Compared to the «Baseline» treatment of perfect observability, trust in *both* of the «Monitoring» treatments («Costless» and «Costly») was *not* lower. The second result is that the efficiency of economic interactions does not necessarily suffer when trustworthiness can only be ascertained through costly monitoring. Joint payoffs in the «Costly Monitoring» treatment were as high as those in the «Costless Monitoring» and the «Baseline» condition, even after taking monitoring expenses into account, though the distribution shifts in favor of trustees when monitoring is costly. Our third result is that trust can be supported by a limited

---

[8]For Elster (2007), it is the latter what truly means to trust: he argues that trust is «the result of two successive decisions: *to engage in the interaction and to abstain from monitoring the interaction partner*» (p. 345, emphasis added). The standard trust game does not include the second step. The monitoring game presented here, however, allows exactly for this succession.

[9]This implementation is very similar to the theoretical model of Miyagawa et al. (2008). See section 3 for a justification of this particular design choice.

[10]There is little guidance on the behavioral implications of a monitoring option. Using monitoring could be emotionally costly to trustors if they have an innate preference for being in a trust relationship (Elster, 2007). Relative to the «Baseline» treatment, the availability of the option could also prime trustors towards distrust and, hence, reduce trust (Burnham et al., 2000).

amount of information on trustworthiness. Similar levels of trust across our three treatments were supported by different information structures, in particular when monitoring was costly: In the «Costless Monitoring» condition, trustors chose to monitor every single action by the trustees. In the «Costly Monitoring» condition, only about half of the trustees' actions were monitored. In most rounds, therefore, trustors chose not to monitor when this was costly, with a consistent dynamic pattern of frequent monitoring in early rounds and sporadic inspections in later rounds. Taken together, these results provide first experimental evidence that the maintenance of trust not only imposes minimal requirements on the structure of repeat interactions, but also smaller requirements on the observability of trustworthiness than expected. In addition, the dynamics of information acquisition provide a window onto how trustors choose to learn about the co-player's trustworthiness. This has interesting parallels to neuroeconomic evidence on how trustors build mental models of their trustee as their reputation develops (King-Casas et al., 2005).

In the remainder we proceed as follows. After a review of the related experimental literature in section 2, we describe the experimental design, procedures, and implementation in section 3. The results are presented in section 4. We summarize and conclude in section 5.

2 Related experimental literature

The present study lies at the intersection of three experimental literatures. The first is a body of research that examines the impact of exogenously manipulating or restricting information within the standard repeated trust game. Burnham et al. (2000) examine the role of priming trustors by introducing the co-player as a «friend» or «foe» and find that despite the priming, learning dynamics lead to behavioral convergence of the two conditions. Anderhub et al. (2002) study a finite-horizon trust game in which trustors are imperfectly informed about the *type* of co-player (completely trustworthy or opportunistic) they are matched with. They find that the aggregate dynamics of this modified trust game match predictions of the reputation formation hypothesis in repeated games. Our design also uses a modified finite-horizon trust game with fixed matches, but differs from both papers in that the information structure within the fixed interaction evolves endogenously according to the choice of the trustor.

The second literature examines the level and evolution of trust when agents take repeat decisions, but not necessarily in fixed matchings. Against this background, different designs examine the impact of exogenous variations in the information structure on trust. Bohnet & Huck (2004) compare fixed matching and random rematching with and without providing to trustors the *history* of the trustee's actions in previous interactions with third parties. They find that a repeat interaction environment with random re-matching and information about their co-players' history is essentially as efficient as a fixed matching environment. Bracht & Feltovich (2009) add a pre-stage to the standard trust game in which the trustor in a mutually one-shot

interaction receives either information about the trustee's *last decision* with another trustor or cheap talk from the current trustee or both. They find that the observability of the recent history of actions leads to high levels of cooperation. Charness et al. (2011) examine the role of indirect reciprocity in a random-rematching trust game and compare the effect of providing information on the co-player's *history* of actions in different *roles*, as a trustor and a trustee. They find strong evidence that information on past behavior as a trustor is as effective for reputation building as past behavior as a trustee, affirming the role of indirect reciprocity as a mechanism supporting trust. The general conclusion from this literature is that information-rich environments are highly conducive towards trust and trustworthiness.[11] Huck et al. (2012) confirm this conclusion, but also find endogenous matching (competition) outperforms exogenous rematching, even when the information environment is exogenously restricted to trustee's identity rather than their history of play. Our paper differs in two important aspects from this literature: Our subjects interact repeatedly in stable pairs across all rounds (no rematching) and, more importantly, the specific information structures are endogenously determined by the trustor rather than being exogenously imposed.[12] These differences reflect our specific interest in how trust responds to poor, but remediable information environments.

The third literature to which this research relates examines the role of costly monitoring in principal-agent relationships. Nagin et al. (2002) conduct a field experiment on how variations in the probability that their sales figures will be audited impacts on the trustworthiness of online call center employees. They find that rational cheating is the dominant behavior, but also that there is considerable heterogeneity: A significant share of employees do not decrease their trustworthiness in response to a decrease in the monitoring probability. Dickinson & Villeval (2008) study the choice of a costly monitoring intensity by a first-moving principal and the effort response by second-moving agents with two treatments, a stranger or partner matching protocol and a payoff function for the principal that was either increasing in the agent's output or fixed. They find that effort is increasing in monitoring intensity, but that this disciplining effect on opportunism is tempered by a crowding out effect in the partner matching protocol. We differ from this line of research in important ways. One is that

---

[11]Interestingly, in contrast to the literature on other social dilemma games (Holcomb & Nelson, 1997; Sainty, 1999; Aoyagi & Fréchette, 2009; Bornstein & Weisel, 2010; Grechenig et al., 2010; Ambrus & Greiner, 2012; Dreber et al., 2014), we are not aware of experimental results in the repeated trust game in which co-player's actions are observed with a certain error probability.

[12]Within the trust game paradigm, there is also an interesting parallel to a recent paper by McEvily et al. (2012) whose experimental design introduces a costly option of insuring against vulnerability in the investment game. In a trust measurement experiment, they offer second movers in series of five one-shot exchange games with strangers the possibility to change the structure of the interaction through a costly option to avoid being exposed to the first mover's decision. They find that subjects apply the option selectively based on expected trust. Some readers of the present paper have commented that the decision to monitor can also be interpreted as an option that reduces vulnerability relative to a situation in which the trustor does not learn about the trustee's actions. However, if vulnerability is the main concern of trustors, simply choosing not to trust provides full insurance against vulnerability in our design.

in our design, the first mover decides on a round-by-round basis whether to engage in monitoring, rather than setting a monitoring probability for all rounds. This choice reflects our interest in the monitoring behavior across rounds from which we hope to learn something about the demand for information on trustworthiness. Also, in our design, monitoring is private knowledge of the first mover while the monitoring intensity in these two papers is public information. Finally, we include a costless monitoring treatment, which allows us to disentangle two separate dimensions of moving from automatic observation to costly monitoring: The fact that monitoring requires a deliberate decision and the fact that monitoring is typically costly.

3 The Experiment

3.1 Experimental game and design

The stage game of the experiment is the well-known (binary) trust game (see e.g. Camerer & Weigelt, 1988; Anderhub et al., 2002; Engle-Warnick & Slonim, 2004, 2006b,a; Slonim & Guillen, 2010).[13] The trustor chooses between investing (option «pink» in the instructions) and an outside option («yellow»). If the outside option is chosen, both players get 15 tokens and the period ends. If the trustor chooses to invest, the period continues with the trustee's choice between splitting (option «brown») or keeping (option «blue»). If the trustee cooperates, she gets 25 tokens and her co-player 30 tokens. Otherwise, she exploits the trustor by taking 50 tokens for himself while his co-player gets 5 tokens.[14] As a «Baseline» condition (BSL in what follows) that directly replicates previous research, we implemented a standard 12-fold repetition of the stage game with perfect information: Each player was informed about the co-player's action automatically, at no cost, without error, and without delay. This treatment constituted the benchmark of the experiment.

The two treatment conditions of the experiment introduced imperfect information into the trust game. In both treatments, a cooperating trustor was no longer automatically informed about the trustee's action. Specifically, without knowing the trustee's action, a trustor decided whether she wanted to monitor the trustee's action or not. If so, the trustor was informed about whether their co-player responded with «brown» or «blue», respectively, at the end of the round. Otherwise, she received no information. Trustees were never informed about whether their co-player monitored them or not.[15] In the first treatment, the «Costless Monitoring» treatment (CSM in what

---

[13]The trust game prototypically captures situations in which efficiency enhancing cooperation is threatened by a possibility of unilateral exploitation. The sequential structure and its simplicity renders it easy for subjects to understand and the interpretation of observed behavior is less difficult than in simultaneous-move games.

[14]This parametrization is rather standard and intended to generate a fair amount of variance in the data: There is an attractive gain from cooperation, but also a quite lucrative incentive for second movers to cheat. Furthermore, the mutual cooperation payoffs are intentionally asymmetric in order to avoid a «fair focal point» (Bohnet et al., 2005; Huck et al., 2012).

[15]We implemented «hidden monitoring» for two major and one minor reasons. First of all, it is an empirically accurate representation of many real-world interactions. Monitoring activities are rarely

follows), monitoring the trustee required a deliberate decision to do so, but involved no cost. In the second treatment, the «Costly Monitoring» treatment (CYM in what follows), trustors had to incur a fee of five tokens in order to acquire this information. Except for these variation, both treatment conditions and the benchmark were exactly identical. We used a between-subjects design to assign treatments to participants.

In each round, a trustor's degree of confidence that the trustee will behave trustworthy is a key variable for her decision to trust or not. Likewise, a trustee's belief in being monitored in a given round is important for her decision to behave trustworthy or opportunistic. In order to learn something about those beliefs, their dynamics, and their relation to observed behavior, we supplemented the experimental game by (non-incentivized) elicitation of the participant's first-order beliefs about their co-player's behavior in the current period. In each period, before any decisions were made, trustors were asked to state their belief about whether their co-player will respond with «brown» or «blue» to «pink», and trustees were asked to state their belief whether their co-player will play «pink» or «yellow». Given that «pink» was played in our main conditions, trustees were asked after their decision to state their belief that their decision will be monitored.

3.2 Subjects and procedures

Participants were recruited from the general undergraduate student population of the University of Heidelberg using the online recruitment system ORSEE (Greiner, 2004). In total 152 subjects participated of which 52.6 percent were female and 85.5 percent German. The mean age was 23.3 years. Subjects were randomly assigned to treatment conditions, 36 to the baseline condition, 56 to the costless monitoring condition, and 60 to the costly monitoring condition. No subject participated more than once or in more than one treatment condition.

All experiments were conducted at the experimental laboratory of the Alfred-Weber-Institute (AWI-Lab) at the University of Heidelberg. Upon entering the laboratory, subjects were randomly assigned to computer terminals. Besides each terminal, an empty sheet of paper and a pen was prepared which participants were allowed to use for taking notes during the experiment. They were instructed to take this sheet

---

performed continuously in practice but rather take the form of surprise inspections or snap samples. For example, while technical surveillance equipment, such as cameras, is often overtly installed, it is typically not running or not tracked permanently. Thus, in such settings the trustees know of the possibility of being monitored at all times, but not at which points in time they are actually monitored—just as in our experiment.

Second, observable monitoring creates the possibility of complex reciprocity effects (Fehr & List, 2004; Falk & Kosfeld, 2006; Sliwka, 2007; von Siemens, 2013): monitoring may be viewed by the trustee as control or a signal of distrust which is evaluated as an offense, justifying to be opportunistic. Conversely, the trustor may strategically use non-monitoring to signal trust and hence induce trustworthiness in the trustee. While such effects are an interesting avenue for further research, the hidden monitoring setting is simpler to interpret and hence a reasonable starting point.

Finally, the way we implement monitoring closely relates to the theoretical literature on repeated games without communication (Miyagawa et al., 2008).

with them after the experiment to ensure that nobody, including the experimenters, could observe their notes. Booths separated the participants visually, ensuring that they made their decisions anonymously and independently. Direct communication among them was strictly forbidden for the duration of the entire session. Furthermore, subjects did not receive any information on the personal identity of any other participant, neither before nor while nor after the experiment.

At the beginning of the experiment, that is, before any decisions were made, subjects received detailed written instructions that explained the exact structure of the game and the procedural rules (see supplementary material). All subjects received the same instructions (only the monitoring fee being replaced across conditions) and this was commonly known. The experiment was framed in a sterile way using neutral language and avoiding value laden terms in the instructions. Post-experimental debriefings attested that no participant had difficulties in comprehending the instructions. The experiment was programmed and conducted with z-Tree (Fischbacher, 2007).

The exact timing of events was as follows. First, the subjects were randomly matched into groups of two. Then twelve rounds of the experimental game described above were played. The binary decisions were made by input boxes to be marked with the computer mouse, beliefs were indicated by a screen slider with a resolution of 100 points. After the twelve rounds, subjects were asked to answer a short questionnaire while the experimenter prepared the payoffs. Subjects were then informed about their payoffs, and then individually called to the experimenter booth, paid out (according to a random number matched to their decisions; no personal identities were used throughout the whole experiment) and dismissed.

In every session subjects received a fixed show-up fee of €3, which was not part of their endowment. The average session had a duration of about 40 minutes and subjects earned €11.37 (€0.03 per token earned) on average, including the fixed show-up fee, with a minimum of €6.75 and a maximum of €15.15. Average earnings exceed the local average hourly wage of a typical student job significantly and can hence be considered meaningful to the participants.

3.3 Predictions

For finite horizon games, the theoretical prediction that trust arises in repeat interactions hinges on the assumption that trustors hold a prior belief that some trustees are committed to reward trust even in the terminal period (Kreps et al., 1982), a belief that is indeed justified as demonstrated by a substantive amount of recent evidence on cooperative behavior (e.g. Henrich et al., 2004; Gintis et al., 2005; Fehr & Schmidt, 2006; Fehr, 2009b; Fischbacher & Gächter, 2010; Thöni et al., 2012). Maintaining a reputation for trustworthiness requires that trustors initiate the trust relationship and observe the trustee's response. The cycle of trust-reward-observation-reputation-trust typically maintains trust in finite-horizon repeat trust games over extended periods before it declines in the final two rounds (Anderhub et al., 2002; Engle-Warnick &

Slonim, 2004, 2006a). We therefore expect to replicate this result in the BSL treatment in which observability of trustworthiness is perfect.

In the presence of limited observability and costless monitoring (the CSM treatment), the only difference to the BSL treatment consists in the trustor having to take an active decision to monitor in order to observe trustworthiness. In a laboratory environment, this involves minimal effort. From a (standard) theoretical perspective, there is therefore no reason to expect a difference in behavior between the CSM and the BSL treatment. Behaviorally, using monitoring could be emotionally costly to trustors if they have an innate preference for being in a trusting relationship: Monitoring «might be incompatible with the agent's emotional attitude toward the other person» (Elster, 2007, p. 346), because she views it as an admission of distrust towards the trustee (McEvily et al., 2012). Relative to the Baseline treatment, the availability of the option could also prime trustors towards distrust and, hence, reduce trust (Burnham et al., 2000). In light of the limited evidence on the presence and scale of these putative mechanism, however, we predict no measurable impact of costless monitoring on trust.

In the presence of limited observability and costly monitoring (the CYM treatment), there are several reasons for expecting that trust will be lower in the monitoring game.[16] One is that every instance of monitoring reduces the gains from trusting: The maximum gain from trust in a given round is 15 tokens while the cost of monitoring is 5 tokens. The second is that trustors who trust and then save on monitoring costs are easily exploited, possibly over multiple periods, without being able to condition future trust on observed trustworthiness. Thirdly, trustees that correctly anticipate non-monitoring in the current round have a higher probability of getting away with cheating, which reduces the relative gains from reputation building. Fourth, repeated trustworthy behavior by the trustee does not accumulate into a reputation unless the trustor incurs the monitoring cost. If the trustee anticipates less observation of trustworthiness due to monitoring costs, the rewards from investing in reputation are reduced relative to perfect observability. We therefore predict that both trustors will respond to the costly monitoring environment with lower trust, and

---

[16]There are three «folk theorems» for infinite discounted games with costly observation, constructing sequential equilibria that can support any (or almost any) outcome, independently from the level of observation costs (Ben-Porath & Kahneman, 2003; Miyagawa et al., 2008; Awaya, 2014). Generally, «folk theorems» are possibility results and of limited utility in deriving experimental predictions, because they do not «predict» anything sharper than that (almost) everything is possible under the restrictions imposed. In addition, the restrictions are demanding: Under reasonable behavioral (e.g. patience, strategic sophistication) and structural (e.g. coordination devices) assumptions the equilibria appear difficult to attain (Gintis, 2009). However, apart from these general issues, those results actually back our predictions since our experimental design rules out the equilibria constructed in the above papers: First, Ben-Porath & Kahneman (2003) and Awaya (2014) require explicit communication which is impossible in our experiment. Second, Miyagawa et al. (2008) requires a public randomization device, which is absent in our experiment. Third, and most importantly, all require infinite or indefinite repetition; under finite repetition, as in our experiment, the equilibria break down. In addition, Awaya (2014) considers random matching games, and is therefore *per se* not applicable to our experimental setup.

trustees with less trustworthiness.

4 Results

We proceed with the presentation of results as follows. In a preliminary step, we conduct a basic replication check that the BSL condition can serve as a benchmark. We then compare behavior across all three treatment conditions.

4.1 Replication check

Both in terms of levels and in terms of dynamics, the results from the BSL condition closely correspond to previously published evidence on trustor and trustee behavior in the finite horizon trust game (Anderhub et al., 2002; Engle-Warnick & Slonim, 2004, 2006a). For example, in their elaborate finite horizon trust game experiment in which a supergame runs for five rounds, Engle-Warnick & Slonim (2004) find that trusting behavior decreases strongly across rounds: trustors trusted in around 60-90 percent of the time in the first round, but only about 10-25 percent of the time in the final round. The lions share of this drop happens in the final two periods. Between 80 and over 90 percent of the trustees respond trustworthy in the first period, and this share exhibits a similar declining trend.[17] Interestingly, these patterns repeat even after subjects have gained a significant amount of experience: trustors still start trusting in the first period in over 80 percent of the time after 30 and more repetitions of the supergame (with different co-players).[18]

In the BSL condition of our experiment, 77.8 percent (14 out of 18) of the trustors trusted in the first, and 22.2 (4 out of 18) in the terminal round—in between the typical pattern of frequent cooperation (about two in three cases) until close to the end and a sharp decline in the final two periods emerges. The full pattern is depicted in figure 2a. Pooled over all rounds, trustors trusted in 63.9 percent (138 out of 216) of the time, and trustees were trustworthy in 82.6 percent (114 out of 138) of the time. The average (per period) joint payoff was 46.0 tokens, 21.8 for trustors and 24.2 for trustees, such that 63.9 percent (16.0 out of 25) of the gains from cooperation were realized. Taken together, those results fit nicely into the range of previously presented evidence on repeated trust games. On this basis, we progress to the key results of the experiment in which we investigate the treatment effects relative to a baseline that is in line with the literature.

---

[17]The pattern depends on whether the share is calculated relative to all trustees or relative to the number of trustees who have been trusted. In the latter case, the decline is generally less steep and often non-monotonic because there is a selection effect over time: those trustees who have not been trustworthy in early periods are simply not trusted any more later on.

[18]These patterns have also been documented in numerous experiments using other social dilemma games (see Clark & Sefton, 2001, Dal Bó, 2005, Andreoni & Croson, 2008, Fischbacher & Gächter, 2010, and the references therein).

Figure 1: Behavior and payoffs by treatment condition, pooled over all rounds. The height of the solid bars indicates the share of trusting actions with (dark) and without (light) monitoring. The dropline indicates the share of trustworthy actions, drawn relative to the height of the bars, respectively, such that the dropline has the same height as the bar if all trustees behaved trustworthy. The outlined bars indicate the joint payoff per group and round, including the cost of monitoring.

4.2 Treatment effects

Here we compare levels of trust, joint payoffs for trustor and trustee, and the information structures between the BSL, the CSM and the CYM treatment, expecting to find no difference between the first two and significantly lower trust and efficiency in the last treatment. Carrying out the comparison leads to three key findings, which are illustrated at a single glance in figure 1 for inspection. All statistical tests reported in this section are Mann-Whitney rank-sum tests (for pairwise comparisons) or Kruskal-Wallis equality-of-populations rank tests (for comparisons over all three conditions) applied to a cross-section in which each observation is a group-level average taken over all twelve rounds.[19]

The first finding is that trust is not lower when trustworthiness can only be observed through active monitoring, irrespective of whether the cost is zero or positive. In the BSL treatment, first movers trusted 63.9 percent of the time (see section 4.1) and in the CSM condition, 69.4 percent of the time (233 out of 336 cases). The difference is not statistically significant ($p = .7411$). In the CYM condition, they trusted no less often (254 out of 360 cases, or 70.6 percent) than in the CSM condition ($p = .5153$) or the BSL condition ($p = .5833$). Overall, the frequency of trust is not significantly different across all three conditions ($p = .7560$). This evidence therefore

---

[19]Note that the individual observations in our data set are not independent in a rigorous statistical sense, that is, strictly speaking each of the 76 matches constitute one independent observation. The procedure used here follows Vanberg (2009), and takes account of this fact.

confirms one part of our predictions: Moving from automatic observation of trustworthiness to deliberate monitoring at zero cost does not significantly affect trust. However, the evidence refutes another part of our prediction, that there is less trust under costly monitoring: Despite a less favorable information environment, there was no less trust if trustworthiness was costly to monitor than under perfect observability.

The second finding, also illustrated in figure 1, is that the efficiency of economic interactions does not necessarily suffer when trustworthiness can only be ascertained through costly monitoring. Average joint payoffs, *including* monitoring costs, were *not* significantly lower in the CYM treatment (46.0 tokens) than in the CSM condition (47.3 tokens, $p = .3409$) or the BSL condition (46.0 tokens, $p = .6012$). Overall, the joint payoffs are not significantly different across all three conditions ($p = .6579$). While expenditures on monitoring were, of course, zero in the BSL and the CSM conditions, the average trustor spent 19.8 tokens per match on monitoring in the CYM condition, after all 6.8 percent of her gross payoff (290.3 tokens). Given the finding on trust levels, this rules out the possibility that continuous monitoring explains the invariance of trust to monitoring costs. Continuous monitoring would have absorbed up to one third of the gains from cooperation, leading to lower overall efficiency. Reconciling these two results requires that the treatment variations induced different monitoring behavior without affecting trust.

The similarities and differences in monitoring behavior across treatments lead to the third finding. Comparing the trustors' information structure in the BSL treatment and the CSM treatment, we find that they are identical: Not a single one among the 233 instances of trust in the CSM condition went unmonitored. Trustors were therefore *de facto* perfectly informed about the entire history of the game at any time, just as in the BSL condition. The introduction of an active monitoring step alone therefore not only leaves the behavioral trust measure unchanged, it also fails to induce any «trust without verification». Comparing the CYM treatment with the CSM treatment, we find a significant response to cost. Trusting first movers monitored less than half of the time (119 out of 254 cases, or 46.9 percent) and trusted without monitoring in more than half of the interactions. This monitoring frequency differs significantly from the CSM condition ($p = .0000$). While trustors' information structures were *exogenously* perfect in the BSL treatment and *endogenously* perfect in the CSM treatment, they were *endogenously imperfect* in the CYM treatment.

Summarizing the three key results, we find that an environment in which trustworthiness is automatically observed and one in which trustworthiness needs to be deliberately monitored do not necessarily differ in their behavioral trust measures or efficiency, even when monitoring involves a non-negligible cost. The information structures that support those identical levels of trust, however, do differ. When trustworthiness can only be observed at a cost, trustors acquire only a fraction of the available information on trustworthiness and yet maintain the same level of trust.

The discrepancy between the predictions and the experimental results raises the question what the dynamic patterns of trust, trustworthiness, and information acquisition are that enable trust to be maintained under costly monitoring. The following

section moves beyond the aggregate comparison across rounds to look at these dynamics in greater detail.

4.3 Dynamics of Monitoring and Trust

Figure 2 shows, for each treatment and round, the share of trustors that trusted with and without monitoring and the share of trustworthy trustees. We first compare the BSL condition (figure 2a) used as a replication check and the CSM condition (figure 2b). Given that trustworthiness is perfectly observed in both treatments, the evolution of trust shows very similar patterns, thus yielding little additional insight beyond a comparison aggregated across rounds as in section 4.2. Comparing this pattern to the pattern in the CYM condition (figure 2c), we find a similar dynamic pattern in terms of total trust, but a clear shift from trust with monitoring towards trust without monitoring over time. In the first period, the vast majority of trustors incurred the monitoring cost in order to observe the trustworthiness of their co-player. Over time, however, the share of monitoring trustors becomes successively less frequent.

Figure 3 illustrates the intertemporal shift towards «trust without verification» more clearly, depicting the frequency of trust with and without monitoring, respectively, as a fraction of all instances of trust. It shows how trustors predominantly resort to monitoring at the beginning of the repeat interaction and then maintain trust without monitoring towards the end. The positive monitoring cost therefore reveals an intertemporal structure of the demand for information on trustworthiness that a deliberate monitoring decision alone does not uncover. The intertemporal pattern of monitoring seen in figure 3 also has an interesting parallel with the neuroeconomic literature that studies how neural responses to each other's decisions evolve in trustor's and trustee's brain over the course of the repeated interaction in a trust game (e.g. King-Casas et al., 2005; Delgado et al., 2005; Krueger et al., 2007; Fareri et al., 2012). King-Casas et al. (2005), in particular, find evidence that trustors engage in «model building of the partner», a process that they find to be complete after five rounds of interaction in a ten-round finite-horizon trust game setting. In our data, we find a similar pattern, in which demand for information on the co-player's action is initially high and then declines such that non-monitoring becomes the modal behavior after round 5.

Figure 4 summarizes information about the belief dynamics of trustors and trustees across rounds. Figure 4a shows the average belief of trustors in the trustworthiness of their co-player. A Kruskal-Wallis test bears out the visual message that those beliefs do not differ between treatments, neither time-averaged ($p = .6855$) nor in any individual round ($p \geq .1306$). Part of the answer why the theoretical prediction fails is therefore found in the fact that the average trustor does not express an expectation that trustees strategically respond to the presence of monitoring costs. Instead, trustors attach significant information value to monitoring in early rounds of the repeat interaction. This suggests that trustors expect to be able to screen among the trustees for behavioral «types» that exhibit some degree of behavioral persistence across rounds. This interpretation is further supported by considering only the class

14

Figure 2: Behavior dynamics by treatment condition. The legend is identical to figure 1.

(a) BSL condition

(b) CSM condition


(c) CYM condition

Figure 3: Frequency of monitoring as a fraction of all instances of trust in the costly monitoring condition.

of trusting trustors—figure 4b replicates figure 4a under this restriction. In the terminal period, those trustors are significantly more confident in the CYM condition than in the CSM condition ($p = .0256$).[20]

The average trustee, on the other hand, does anticipate that trustors monitor less over time in the CYM condition while they expect monitoring to remain constantly high in the CSM condition. This is illustrated in figure 4c. Mann-Whitney tests show that trustees' beliefs in being monitored are not significantly different between treatments in the first round ($p = .4593$), but significantly different in all subsequent rounds ($p = .0323$ and $p = .0687$ in the second and third round, respectively, $p \leq .0256$ in the remaining rounds).

The combination of average trustors that monitor in early rounds, but expect little strategic behavior from trustees, and average trustees that anticipate this pattern explains the impact that monitoring costs have on the distribution of payoffs. When monitoring is costly, trustees acting strategically receive a greater share of the joint payoffs because they respond to the incentive to cheat in the later half of a match in the CYM condition because late cheats have a higher likelihood of remaining undetected and hence without impact on reputation.[21]

Figure 5 shows the realized payoffs over time in all three treatment conditions.

---

[20]The difference between CYM and BSL is not statistically significant ($p = .1563$), but note that there are only four observations in the BSL (compared to 8 and 11 in the CSM and CYM conditions, respectively.

[21]Conventional economic theory suggests that the temptation to cheat is decreasing in the perceived likelihood of being detected, and *vice versa*. On average, trustees respond consistently with this prediction: In both the CSM and the CYM conditions, the average cooperating trustee had a stronger belief of being monitored (.920 in CSM, .680 in CYM) than the average defecting trustee (.887 in CSM, .467 in CYM). Trustworthy behavior and the belief of being monitored is positively correlated in the CYM condition (Kendall's $\tau_b = 0.206$, $p = .0002$), but not in the CSM condition ($\tau_b = 0.043$, $p = .4859$).

16

Figure 4: Dynamics of beliefs by treatment. Panel (a) depicts the average trustor's belief in their trustee acting trustworthy, panel (b) depicts the same but restricted to the actually trusting trustors. Panel (c) depicts the average trustee's belief in being monitored.
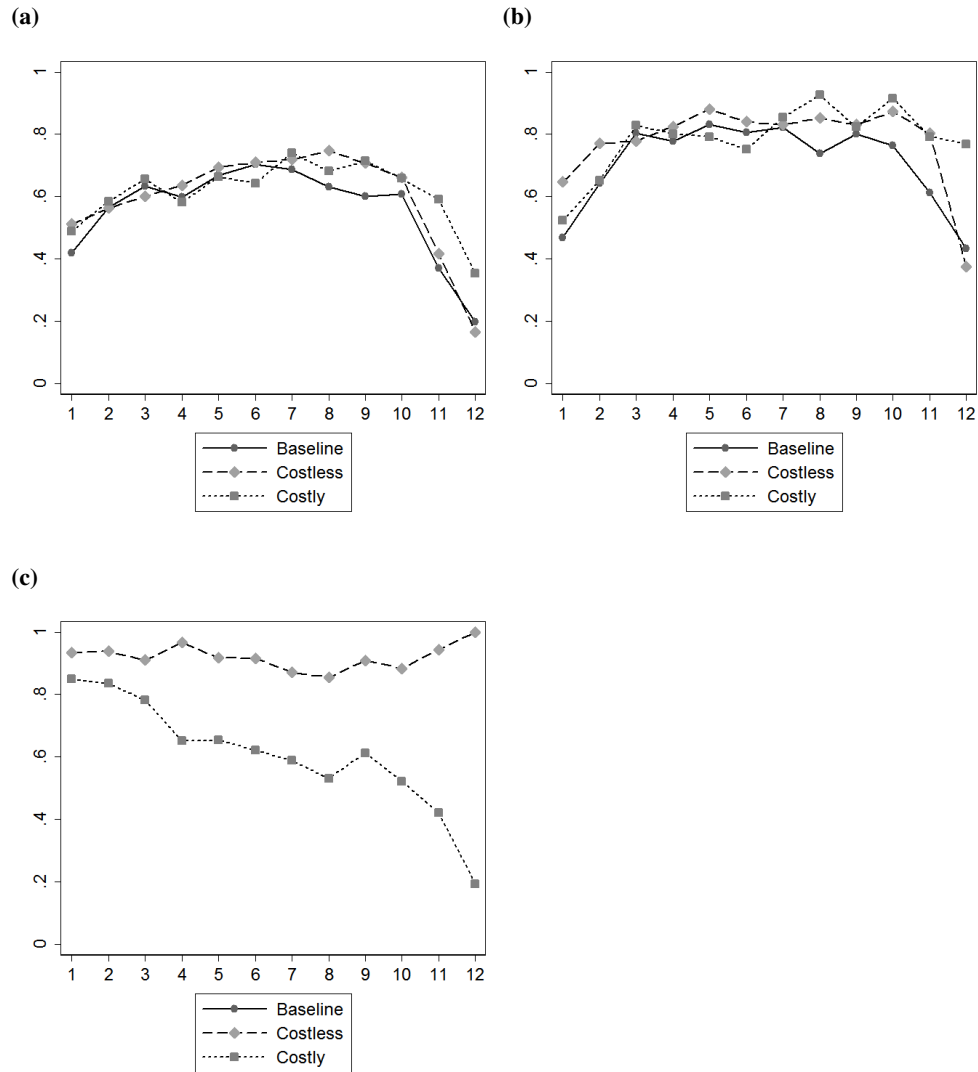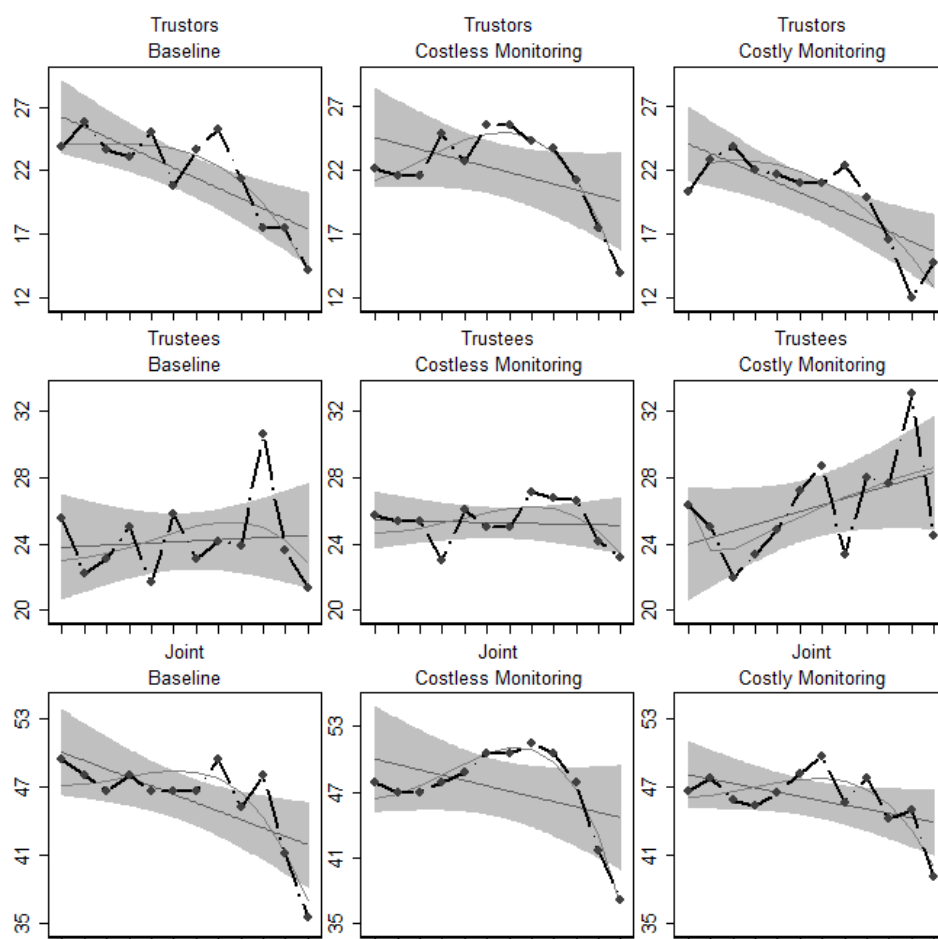
(a)

(b)

(c)

Figure 5: Average payoffs over time. The connected dots depict the averages per round. They are supplemented by a linear fit with 95-percent confidence interval (shaded area) and fractional polynomial fit.

In the BSL condition, trustors reaped on average 47.2 percent of the joint payoff (minimally 39.8 percent in round 12, maximally 53.8 percent in round 2), in the CSM condition on average 46.3 percent (minimally 37.5 percent in round 12, and maximally 51.9 percent in round 4). In the CYM condition, the average trustee reaped substantially larger payoffs in the second half of the match, both compared to the first half and the other two treatment conditions. In the first six periods, the average trustor received on average a share of 47.0 percent of the joint payoff (49.8 percent in the BSL and 47.9 percent in the CSM condition), in the final six periods 39.0 percent (44.7 in the BSL and 44.8 in the CSM condition) with a minimum in the penultimate period (26.7 percent). A significant number of trustees therefore acted on the incentive to cheat in the second half of the match.

An examination of the dynamics of trust at the individual subject level substantiates the presence of significant heterogeneities behind the average results and add some additional subtlety, in particular with respect to the beliefs underpinning individual behavior (see appendix A). Individual-level data in the CYM condition show that (i) most trustors consistently gathered information in early rounds of the repeat interaction even when their beliefs about the trustee's trustworthiness were initially pessimistic; (ii) monitoring was reduced only if the trustee was observed to be trustworthy early in the interaction; and (iii) among a significant subset of trustors, the trust and monitoring dynamics exhibit large heterogeneity that is not fully explained by beliefs. Examining the elicited beliefs more closely, we find that trust without monitoring is positively related to the trustor's confidence in their co-player's trustworthiness.[22] Using post-experimental survey data on individual preferences, we also find that the incidence of trust without monitoring is negatively related to the degree of risk aversion and betrayal aversion (see appendix B).

5 Conclusion

Recent experimental evidence based on the trust game has revealed how surprisingly robust trust is to structural imperfections in the environment in which the economic interactions take place. Adequately rich information environments are sufficient to ensure that trust survives in repeat interactions with as little structure as random rematching. This raises the question of whether trust is perhaps much more vulnerable to imperfections that compromise the richness of the information structure. One common and much discussed imperfection is the one that we studied in this paper: trustors in many relevant contexts need to engage in (costly) monitoring in order to ascertain whether the trustee rewarded their trust or not.

---

[22]Correlation between the trustors' belief about their trustee acting trustworthy and their own trusting behavior is strongly positive and significant in the BSL condition, (Kendall's τ = .444, p = .0000), the CSM condition (τ = .627, p = .0000) and the CYM condition (τ = .558, p = .0000). However, separating in the CYM condition trust with monitoring and trust without, the latter turns out to be correlated more strongly with trustors' beliefs in trustworthiness (τ = .378, p = .0000) than the former (τ = .152, p = .0008).

Compared to a setting with perfect observability, a setting with limited, but remediable observability of trustworthiness is predicted to give rise to less efficient interactions. The reason is that the effort and cost of monitoring negatively affects both the relative gains from trust versus opting out for the trustor and the relative gains from reputation-building versus immediate exploitation for the trustee. This means that when the limited observability of trustworthiness can only be remedied at a cost, trust should either be lower or the aggregate costs of adequately monitoring trustworthiness should reduce the joint surplus. Despite these factors, neither the level of trust nor efficiency was negatively affected in the repeated interactions in which the observability of trustworthiness was compromised. The explanation is that when trustees are mostly trustworthy, trustors demand relatively little information on trustee's trustworthiness in order to build up and maintain trust. The average trustor ostensibly believes in behavioral «types», i.e. some degree of behavioral persistence in trustee's actions, and does not express an expectation of sophisticated Bayesian strategizing on behalf of the trustee. Instead, they rely on later spot checks to confirm that their identification strategy in early rounds was correct. As a result, the economic environment with an information imperfection performed as efficiently as the baseline treatment with perfect information, even though the distribution of gains differed in favor of the trustees. Putting these results in the context of the larger literature, they point to the possibility that the maintenance of trust not only requires minimal structure in the repeat interactions, but also considerably less observability of trustworthiness than previously thought. We commented on the parallels to neuroeconomic studies on the evolution of trust above.

We see a number of promising avenues of further research on the basis of variations of the monitoring game presented here. Real-world settings differ in monitoring costs, and even while we consider the costs in the present experiment (at one-third of the per-round gain from rewarded trust) as large and meaningful, there is no natural upper limit to monitoring costs. This raises both the question of the shape of trustors' demand curve for information about trustworthiness and the question of the existence and level of a possible threshold of monitoring costs for trust to break down irreparably. The monitoring game also excluded deliberately the question of the observability of trustor's monitoring actions by the trustee. While the experimental literature on monitoring reviewed in section 2 goes some way towards understanding the higher-level game that can arise through trustors sending signals of high trust by observably not monitoring, the stochastic nature of the monitoring rate in a multi-agent setting explored there fails to exhaust the full ramifications of observable non-monitoring in repeat interactions, in particular in fixed matches (Elster, 2007).

An interesting avenue for further research is a detailed investigation of the strategies the players play in the monitoring game and their equilibrium properties. Specifically, the expectation that there is a higher probability of being monitored during early stages may give trustees an incentive to a «higher order» of reputation building: In standard repeated games with perfect monitoring (but incomplete information), strategically acting trustees can maintain a favorable reputation in order to induce

20

the trustor to cooperate until close to termination (Kreps et al., 1982). In the monitoring game, trustees can induce the trustors not only to trust, but also to refrain from monitoring. We suspect that the incentive for the latter («second-order reputation building») is stronger than the former («first-order reputation building») because under perfect monitoring the maximum number of periods in which a strategically acting trustee can exploit the trustor is equal to one (assuming that the trustor will not trust again once cheated), while in the monitoring game there is the possibility of cheating over *multiple* periods once the trustor starts to trust without monitoring. Intuitively, (some) trustees may deliberately try to «earn» a reputation in the initial periods in which they are likely to be monitored, favorable enough to be trusted without verification later on. But this strategy can only be part of some kind of *mixed* strategy equilibrium, because trust without verification with certainty is not a best response to it. The sporadic inspections many trustors perform in later periods is a hint in this direction.

References

Alchian, A. A. & Demsetz, H. (1972). Production, information costs, and economic organization. *American Economic Review*, 62(5), 777–795.

Ambrus, A. & Greiner, B. (2012). Imperfect public monitoring with costly punishment–an experimental study. *American Economic Review*, 102(7), 3317–3332.

Anderhub, V., Engelmann, D., & Güth, W. (2002). An experimental study of the repeated trust game with incomplete information. *Journal of Economic Behavior & Organization*, 48(2), 197–216.

Andreoni, J. & Croson, R. (2008). Partners versus strangers: Random rematching in public goods experiments. In C. R. Plott & V. L. Smith (Eds.), *Handbook of Experimental Economics Results* (pp. 776–783). Amsterdam, The Netherlands: North-Holland.

Andreoni, J. & Miller, J. H. (1993). Rational cooperation in the finitely repeated Prisoner's Dilemma: Experimental evidence. *Economic Journal*, 103(418), 570–585.

Annen, K. (2013). Social capital as a substitute for formality: Evidence from Bolivia. *European Journal of Political Economy*, 31, 82–92.

Aoyagi, M. & Fréchette, G. (2009). Collusion as public monitoring becomes noisy: Experimental evidence. *Journal of Economic Theory*, 144(3), 1135–1165.

Armendáriz, B. & Morduch, J. (2005). *The Economics of Microfinance*. Cambridge, MA, USA: MIT Press.

Arrow, K. J. (1972). Gifts and exchanges. *Philosophy & Public Affairs*, 1(4), 343–362.

Ashraf, N., Bohnet, I., & Piankov, N. (2006). Decomposing trust and trustworthiness. *Experimental Economics*, 9(3), 193–208.

Awaya, Y. (2014). Community enforcement with observation costs. *Journal of Economic Theory*, 154, 173–186.

Ben-Ner, A. & Halldorsson, F. (2010). Trusting and trustworthiness: What are they, how to measure them, and what affects them. *Journal of Economic Psychology*, 31(1), 64–79.

Ben-Porath, E. & Kahneman, M. (2003). Communication in repeated games with costly monitoring. *Games and Economic Behavior*, 44(2), 227–250.

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1), 122–142.

Berggren, N. & Jordahl, H. (2006). Free to trust: Economic freedom and social capital. *Kyklos*, 59(2), 141–169.

Bjørnskov, C. (2011). Combating corruption: On the interplay between institutional quality and social trust. *Journal of Law and Economics*, 54(1), 135–159.

Bohnet, I., Frey, B. S., & Huck, S. (2001). More order with less law: On contract enforcement, trust, and crowding. *American Political Science Review*, 95(1), 131–144.

Bohnet, I., Harmgart, H., Huck, S., & Tyran, J.-R. (2005). Learning trust. *Journal of the European Economic Association*, 3(2-3), 322–329.

Bohnet, I. & Huck, S. (2004). Repetition and reputation: Implications for trust and trustworthiness when institutions change. *American Economic Review*, 94(2), 362–366.

Bornstein, G. & Weisel, O. (2010). Punishment, cooperation, and cheater detection in "noisy" social exchange. *Games*, 1(1), 18–33.

Bowles, S. & Gintis, H. (2002). Social capital and community governance. *Economic Journal*, 112(483), F419–F436.

Bracht, J. & Feltovich, N. (2009). Whatever you say, your reputation precedes you: Observation and cheap talk in the trust game. *Journal of Public Economics*, 93(9-10), 1036–1044.

Broix, C. & Posner, D. N. (1998). Social capital: Explaining its origins and effects on government performance. *British Journal of Political Science*, 28(4), 686–693.

Burnham, T., McCabe, K., & Smith, V. L. (2000). Friend-or-foe intentionality priming in an extensive form trust game. *Journal of Economic Behavior & Organization*, 43(1), 57–73.

Camera, G. & Casari, M. (2009). Cooperation among strangers under the shadow of the future. *American Economic Review*, 99(3), 979–1005.

Camera, G., Casari, M., & Bigoni, M. (2012). Cooperative strategies in anonymous economies: An experiment. *Games and Economic Behavior*, 75(2), 570–586.

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

Carpenter, J. P., Daniere, A. G., & Takahashi, L. M. (2004). Cooperation, trust, and social capital in southeast asian urban slums. *Journal of Economic Behavior & Organization*, 55(4), 533–551.

Cason, T. N. & Khan, F. U. (1999). A laboratory study of voluntary public goods provision with imperfect monitoring and communication. *Journal of Development Economics*, 58(2), 533–552.

Charness, G., Du, N., & Yang, C.-L. (2011). Trust and trustworthiness reputations in an investment game. *Games and Economic Behavior*, 72(2), 361–375.

Clark, K. & Sefton, M. (2001). The sequential prisoner's dilemma: Evidence on reciprocation. *Economic Journal*, 111(468), 51–68.

Cochard, F., Van, P. N., & Willinger, M. (2004). Trusting behavior in a repeated investment game. *Journal of Economic Behavior & Organization*, 55(1), 31–44.

Cooper, R., DeJong, D. V., Forsythe, R., & Ross, T. W. (1996). Cooperation without reputation: Experimental evidence from Prisoner's Dilemma games. *Games and Economic Behavior*, 12(2), 187–218.

Cox, J. C. (2004). How to identify trust and reciprocity. *Games and Economic Behavior*, 46(2), 260–281.

Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review*, 95(5), 1591–1604.

Dal Bó, P. & Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1), 411–429.

Delgado, M. R., Frank, R. H., & Phelps, E. A. (2005). Perceptions of moral character modulate the neural systems of reward during the trust game. *Nature Neuroscience*, 8, 1611–1618.

Dickinson, D. & Villeval, M.-C. (2008). Does monitoring decrease work effort? the complementarity between agency and crowding-out theories. *Games and Economic Behavior*, 63(1), 56–76.

DiPasquale, D. & Glaeser, E. L. (1999). Incentives and social capital: Are homeowners better citizens? *Journal of Urban Economics*, 45(2), 354–384.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., & Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3), 522–550.

Dong, X.-Y. & Dow, G. K. (1993). Monitoring costs in Chinese agricultural teams. *Journal of Political Economy*, 101(3), 539–553.

Dreber, A., Fudenberg, D., & Rand, D. G. (2014). Who cooperates in repeated games: The role of altruism, inequity aversion, and demographics. *Journal of Economic Behavior & Organization*, 98, 41–55.

Duffy, J. & Ochs, J. (2009). Cooperative behavior and the frequency of interaction. *Games and Economic Behavior*, 66(2), 785–812.

Dulleck, U. & Kerschbamer, R. (2006). On doctors, mechanics, and computer specialists: The economics of credence goods. *Journal of Economic Literature*, 44(1), 5–42.

Eckel, C. C. & Wilson, R. K. (2004). Is trust a risky decision? *Journal of Economic Behavior & Organization*, 55(4), 447–465.

Elster, J. (2007). *Explaining Social Behavior: More Nuts and Bolts for the Social Sciences*. Cambridge, UK: Cambridge University Press.

Emons, W. (1997). Credence goods and fraudulent experts. *RAND Journal of Economics*, 28(1), 107–119.

Engle-Warnick, J. & Slonim, R. L. (2004). The evolution of strategies in a repeated trust game. *Journal of Economic Behavior & Organization*, 55(4), 553–573.

Engle-Warnick, J. & Slonim, R. L. (2006a). Inferring repeated-game strategies from actions: evidence from trust game experiments. *Economic Theory*, 28(3), 603–632.

Engle-Warnick, J. & Slonim, R. L. (2006b). Learning to trust in indefinitely repeated games. *Games and Economic Behavior*, 54(1), 95–114.

Falk, A., Gächter, S., & Kovács, J. (1999). Intrinsic motivation and extrinsic incentives in a repeated game with incomplete contracts. *Journal of Economic Psychology*, 20(3), 251–284.

Falk, A. & Kosfeld, M. (2006). The hidden costs of control. *American Economic Review*, 96(5), 1611–1630.

Fareri, D. S., Chang, L. J., & Delgado, Mauricio, R. (2012). Effects of direct social experience on trust decisions and neural reward circuitry. *Frontiers in Neuroscience*, 6, 148.

Fehr, E. (2009a). On the economics and biology of trust. *Journal of the European Economic Association*, 7(2-3), 235–266.

Fehr, E. (2009b). Social preferences and the brain. In P. W. Glimcher, C. F. Camerer, E. Fehr, & R. A. Poldrack (Eds.), *Neuroeconomics: Decision Making and the Brain* (pp. 215–232). London: Academic Press.

Fehr, E., Kirchler, E., Weichbold, A., & Gächter, S. (1998). When social norms overpower competition: Gift exchange in experimental labor markets. *Journal of Labor Economics*, 16(2), 324–351.

Fehr, E., Kirchsteiger, G., & Riedl, A. (1993). Does fairness prevent market clearing? an experimental investigation. *Quarterly Journal of Economics*, 108(2), 437–459.

Fehr, E. & List, J. A. (2004). The hidden costs and returns of incentives–trust and trustworthiness among CEOs. *Journal of the European Economic Association*, 2(5), 743–771.

Fehr, E. & Schmidt, K. M. (2006). The economics of fairness, reciprocity and altruism - experimental evidence and new theories. In S.-C. Kolm & J. M. Ythier (Eds.), *Handbook of the Economics of Giving, Altruism and Reciprocity*, volume 1 of *Handbooks in Economics* (pp. 615–691). Amsterdam, The Netherlands: North-Holland.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.

Fischbacher, U. & Gächter, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review*, 100(1), 541–556.

Fudenberg, D., Levine, D., & Maskin, E. (1994). The Folk Theorem with imperfect public information. *Econometrica*, 62(5), 997–1039.

Fudenberg, D. & Maskin, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3), 533–554.

Fukuyama, F. (1995). *Tust: The Social Virtues and the Creation of Prosperity*. New York, USA: Free

Press.

Gächter, S. & Falk, A. (2002). Reputation and reciprocity: Consequences for the labour relation. *Scandinavian Journal of Economics*, 104(1), 1–26.

Georgarakos, D. & Fürth, S. (2015). Household repayment behavior: The role of social capital and institutional, political, and religious beliefs. *European Journal of Political Economy*, forthcoming.

Gintis, H. (2009). *The Bounds to Reason*. Princeton: Princeton University Press.

Gintis, H., Bowles, S., Boyd, R., & Fehr, E., Eds. (2005). *Moral Sentiments and Material Interests. The Foundations of Cooperation in Economic Life*. Cambridge, USA: MIT Press.

Glaeser, E. L., Laibson, D. I., Scheinkman, J. A., & Soutter, C. L. (2000). Measuring trust. *Quarterly Journal of Economics*, 115(3), 811–846.

Grechenig, K., Nicklisch, A., & Thöni, C. (2010). Punishment despite reasonable doubt–a public goods experiment with sanctions under uncertainty. *Journal of Empirical Legal Studies*, 7(4), 847–867.

Greif, A. (1993). Contract enforceability and economic institutions in early trade: The Maghribi traders' coalition. *American Economic Review*, 83(5), 1281–1302.

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

Guiso, L., Sapienza, P., & Zingales, L. (2004). The role of social capital in financial development. *American Economic Review*, 94(3), 526–556.

Guiso, L., Sapienza, P., & Zingales, L. (2008). Alfred Marshall Lecture: Social capital as good culture. *Journal of the European Economic Association*, 6(2-3), 295–320.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., & Gintis, H., Eds. (2004). *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford, UK: Oxford University Press.

Holcomb, J. H. & Nelson, P. S. (1997). The role of monitoring in duopoly market outcomes. *Journal of Socio-Economics*, 26(1), 79–93.

Houser, D., Schunk, D., & Winter, J. (2010). Distinguishing trust from risk: An anatomy of the investment game. *Journal of Economic Behavior & Organization*, 74(1-2), 72–81.

Huck, S., Lünser, G. K., & Tyran, J.-R. (2012). Competition fosters trust. *Games and Economic Behavior*, 76(1), 195–209.

Jackman, R. W. & Miller, R. A. (1998). Social capital and politics. *Annual Review of Political Science*, 1, 47–73.

Johnson, N. D. & Mislin, A. A. (2011). Trust games: A meta-analysis. *Journal of Economic Psychology*, 32(5), 865–889.

Kandel, E. & Lazear, E. P. (1992). Peer pressure and partnerships. *Journal of Political Economy*, 100(4), 801–817.

Kandori, M. (1992a). Social norms and community enforcement. *Review of Economic Studies*, 59(1), 63–80.

Kandori, M. (1992b). The use of information in repeated games with imperfect monitoring. *Review of Economic Studies*, 59(3), 581–593.

Kanemoto, Y. & MacLeod, W. B. (1991). The theory of contracts and labor practices in Japan and the United States. *Managerial and Decision Economics*, 12(2), 159–170.

Karlan, D. S. (2005). Using experimental economics to measure social capital and predict financial decisions. *American Economic Review*, 95(5), 1688–1699.

King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., & Montague, P. R. (2005). Getting to know you: Reputation and trust in a two-person economic exchange. *Science*, 308(5718), 78–83.

Kirchler, E., Fehr, E., & Evans, R. (1996). Social exchange in the labor market: Reciprocity and trust versus egoistic money maximization. *Journal of Economic Psychology*, 17(3), 313–341.

Knack, S. (2002). Social capital and the quality of government: Evidence from the states. *American Journal of Political Science*, 46(4), 772–785.

Knack, S. & Keefer, P. (1997). Does social capital have an economic payoff? a cross-country investi-

gation. *Quarterly Journal of Economics*, 112(4), 1251–1288.

Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, 27(2), 245–252.

Krueger, F., McCabe, K., Moll, J., Kriegeskorte, Nikolaus, Z. R., Strenziok, M., Heinecke, A., & Grafman, J. (2007). Neural correlates of trust. *Proceedings of the National Academy of Sciences of the United States of America*, 104(50), 20084–20089.

Lazear, E. P. (1993). Labor economics and the psychology of organizations. *Journal of Economic Perspectives*, 5(2), 89–110.

Letki, N. & Evans, G. (2005). Endogenizing social trust: Democratization in east-central Europe. *British Journal of Political Science*, 35(3), 515–529.

McEvily, B., Radzevick, J. R., & Weber, R. A. (2012). Whom do you distrust and how much does it cost? an experiment on the measurement of trust. *Games and Economic Behavior*, 74(1), 285–298.

Miyagawa, E., Miyahara, Y., & Sekiguchi, T. (2008). The folk theorem for repeated games with observation costs. *Journal of Economic Theory*, 139(1), 192–221.

Nagin, D. S., Rebitzer, J. B., Sanders, S., & Taylor, L. J. (2002). Monitoring, motivation, and management: The determinants of opportunistic behavior in a field experiment. *American Economic Review*, 92(4), 850–873.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge, UK: Cambridge University Press.

Ostrom, E. & Gardner, R. (1993). Managing local commons: Theoretical issues in incentive design. *Journal of Economic Perspectives*, 7(4), 113–134.

Ostrom, E. & Walker, J., Eds. (2003). *Trust & Reciprocity: Interdisciplinary Lessons from Experimental Research*. New York: Russell Sage Foundation.

Putnam, R. D. (1993a). *Making democracy work: Civic traditions in modern Italy*. Princeton: Princeton University Press.

Putnam, R. D. (1993b). The prosperous community: Social capital and public life. *American Prospect*, 13(1995), 65–78.

Rousseau, D. M., Sitkin, S. B., Burt, R. S., & Camerer, C. F. (1998). Not so different after all: A cross-discipline view of trust. *Academy of Management Review*, 23(3), 393–404.

Rubinstein, A. (1979). Equilibrium in supergames with the overtaking criterion. *Journal of Economic Theory*, 21(1), 1–9.

Rustagi, D., Engel, S., & Kosfeld, M. (2010). Conditional cooperation and costly monitoring explain success in forest commons management. *Science*, 330, 961–965.

Sainty, B. (1999). Achieving greater cooperation in a noisy prisoner's dilemma: an experimental investigation. *Journal of Economic Behavior & Organization*, 39(4), 421–435.

Sampson, R. J., Raudenbush, S. W., & Earls, F. (1997). Neighborhoods and violent crime: A multilevel study of collective efficacy. *Science*, 277(5328), 918–924.

Sapienza, P., Toldra, A., & Zingales, L. (2007). *Understanding Trust*. NBER Working Paper 13387, National Bureau of Economic Research, Cambridge, USA.

Schechter, L. (2006). Trust, trustworthiness, and risk in rural Paraguay. *Experimental Economics*, 9(2), 173.

Schechter, L. (2007). Traditional trust measurement and the risk confound: An experiment in rural Paraguay. *Journal of Economic Behavior & Organization*, 62(2), 272–292.

Sell, J. & Wilson, R. K. (1991). Levels of information and contributions to public goods. *Social Forces*, 70(1), 107–124.

Shapiro, C. & Stiglitz, J. E. (1984). Equilibrium unemployment as a worker discipline device. *American Economic Review*, 74(3), 433–444.

Sliwka, D. (2007). Trust as a signal of a social norm and the hidden costs of incentive schemes. *American Economic Review*, 97(3), 999–1012.

Slonim, R. & Guillen, P. (2010). Gender selection discrimination: Evidence from a trust game. *Journal of Economic Behavior & Organization*, 76(2), 385–405.

Thöni, C., Tyran, J.-R., & Wengström, E. (2012). Microfoundations of social capital. *Journal of Public Economics*, 96(7-8), 635–643.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46(1), 35–57.

Vanberg, C. (2009). Why do people keep their promises? an experimental test of two explanations. *Econometrica*, 76(6), 1467–1480.

Varian, H. R. (1990). Monitoring agents with other agents. *Journal of Institutional and Theoretical Economics*, 146(1), 153–174.

von Siemens, F. A. (2013). Intention-based reciprocity and the hidden costs of control. *Journal of Economic Behavior & Organization*, 92, 55–65.

Weissing, F. & Ostrom, E. (1991). Irrigation institutions and the games irrigators play: Rule enforcement without guards. In R. Selten (Ed.), *Game Equilibrium Models II* (pp. 188–262). Berlin, Heidelberg: Springer.

Williamson, S. D. (1986). Costly monitoring, financial intermediation, and equilibrium credit rationing. *Journal of Monetary Economics*, 18(2), 159–179.

Williamson, S. D. (1987). Costly monitoring, loan contracts, and equilibrium credit rationing. *Quarterly Journal of Economics*, 102(1), 135–146.

Wolinsky, A. (1995). Competition in markets for credence goods. *Journal of Institutional and Theoretical Economics*, 151(1), 117–131.

Yu, S., Beugelsdijk, S., & de Haan, J. (2015). Trade, trust and the rule of law. *European Journal of Political Economy*, 37, 102–115.

Zak, P. J. & Knack, S. (2001). Trust and growth. *Economic Journal*, 111(470), 295–321.
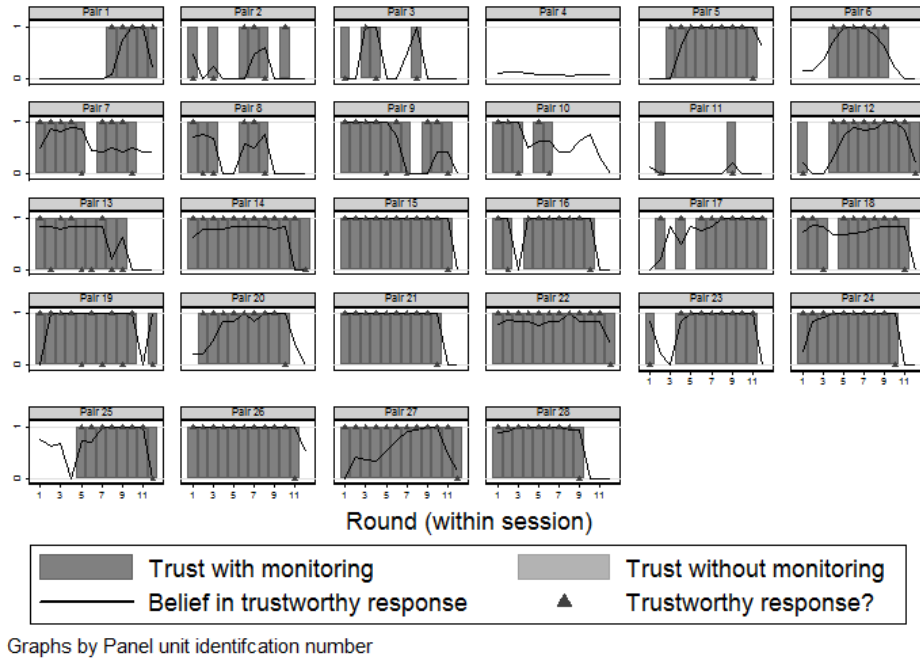
A   Individual-level dynamics

An examination of the individual dynamics of trust reveals interesting heterogeneities behind the averages.

Figure 6a depicts the individual dynamics of trustor behavior in the CSM condition. The bars indicate whether the trustor trusted in a given period, where a bar is shaded in dark gray if accompanied by monitoring and light gray if not. The latter never occurred in the CSM condition. The markers at the top and the bottom of the bars represent the trustee's responses, where a marker at the top means trustworthy behavior and a marker at the bottom means defection. Finally, the black lines depict the trustors' beliefs in a trustworthy response.
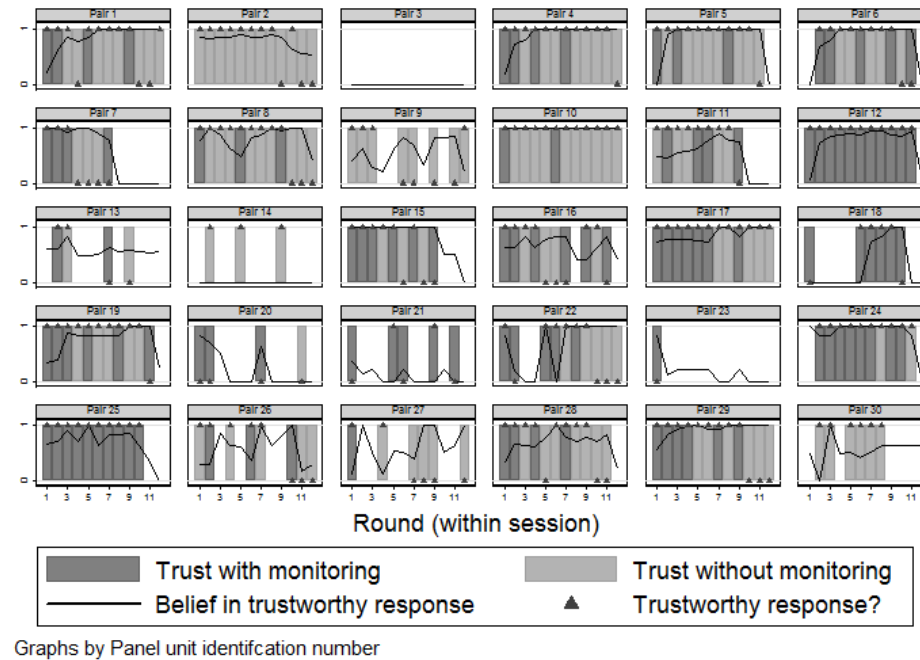
It is evident that dynamics in individual matches differ. Particularly interesting are the individual belief patterns. Almost half of the trustors start with a rather pessimistic prior regarding trustworthiness (see in particular pairs 1–6, 11, 12, 17, 19, 20, 24, and 27). For these subjects, learning requires significant risk-taking. Only one trustor refused to do so (pair 4), and hence forewent all feasible gains from trust, the rest tested the trustee at least once over the course of interaction. Trustors whose trust was not exploited appear to become rapidly more confident. Trustors who «gave it a try» (see pairs 7–10, 13–16, 18, 21–23, 25–26, and 28) and who were disappointed usually became more, or remained, pessimistic and sometimes punished detected cheats (all cheats were detected) in non-terminal periods with at least one period of opting out. Note that almost all trustors anticipated the trustee's strategic incentive to defect in the final period. In sum, this individual investigation reveals that trustor's behavior is broadly consistent with their beliefs about trustworthiness even at the in-

Figure 6: Individual first mover dynamics. Panel (a) depicts the CSM condition, panel (b) the CYM condition.

(a)

Round (within session)

Trust with monitoring | Trust without monitoring
Belief in trustworthy response | ▲ Trustworthy response?

Graphs by Panel unit identifcation number

(b)

Round (within session)

Trust with monitoring | Trust without monitoring
Belief in trustworthy response | ▲ Trustworthy response?

Graphs by Panel unit identifcation number

dividual level, and that their beliefs respond to monitored trustee behavior in expected fashion.

In the CYM condition, the same general conclusion can be drawn, but with the important qualification that the same conclusion arises despite frequent non-monitoring. The majority of trustors started the match with monitoring. If trustees acted trustworthy once or twice, most trustors rapidly moved to trust without monitoring, moderated by occasional monitoring and responding with opting out if monitoring revealed a cheat (see pairs 1, 4–8, 11, 12, 15, 19, 22 and 28–29 for those patterns). Almost half of the trustors behaved in line with a heuristic strategy that can be summarized as shifting from trust with monitoring to trust with only sporadic monitoring as long as no cheating is detected.

Among the other half of trustors, we find considerable individual heterogeneity. There is one first mover (pair 3) who did not trust even once. One first mover (pair 12) started with very pessimistic beliefs, yet trusted the trustee's trustworthiness, but despite becoming very optimistic over time in line with observed high trustworthiness, monitored in every round (spending 55 tokens on monitoring alone). A similar pattern resulted in pair 25. Another trustor (pair 2), starting with a very optimistic prior, trusted without monitoring for all twelve periods, despite foreseeing the trustee's strategic incentive to cheat towards the terminal period. Thus, there clearly appears to be some individual heterogeneity in the propensity to trust and monitor that is not accounted for by beliefs alone.

B Supplementary evidence from post-experimental debriefings

Previous experimental evidence reported negative correlations between risk aversion and trusting behavior in the trust game (Eckel & Wilson, 2004; Houser et al., 2010). Consistent with this, we find trusting behavior in our experiment to be positively correlated with an experimentally validated survey measure of individual risk tolerance. The item contains the question «Are you generally willing to take risks, or do you try to avoid risks?», and respondents answer the question on a 11-point Likert scale ranging from 0 (very risk averse) to 10 (very risk seeking). The item is used in the German Socio-Economic Panel (SOEP) and has been shown to be good predictor of behavior in experiments with decisions under risk (Dohmen et al., 2011). Like in section 4 (and explained there) we use a cross-section in which each observation is a group-level average taken over all twelve rounds, and find a significant and positive correlation between this measure of risk tolerance and trusting behavior in our experiment (Kendall's τ = .197, p = .0196, over all treatment conditions). Restricting the analysis to trust without monitoring in the «Costly Monitoring» condition, the correlation coefficient is of almost identical magnitude, but rejecting the independence hypothesis carries a significant probability of error (τ = .194, p = .1707).

We find interesting correlations of trustor behavior with a measure of betrayal aversion. The measure is based on two items, which are originally part of a battery (also implemented in the SOEP) that measures inclinations for reciprocal behavior:

«If I suffer a serious wrong, I will take revenge as soon as possible, no matter what the cost» and «If somebody offends me, I will offend him/her back». Respondents can answer on a 7-point Likert scale ranging from 1 («does not apply to me at all») and 7 («applies to me perfectly»). It has been argued that the sum of both responses is also good proxy for betrayal aversion (Fehr, 2009a). In our experiment, this measure is uncorrelated with trusting behavior in general ($\tau = .068$, $p = .4219$, over all treatment conditions), *positively* correlated with trust *with* monitoring ($\tau = .234$, $p = .0050$, over all treatment conditions), and *negatively* with trust *without* monitoring ($\tau = -.229$, $p = .0122$, over all treatment conditions). This suggests that betrayal aversion might be a critical trait moderating the monitoring decision: betrayal averse subjects tend to «verify» their co-players trustworthiness after trusting, betrayal tolerant subjects have a stronger tendency trust without «verification».

C   Supplementary experiment with two-round matches

We wondered whether the results found in this paper are robust to a significant shortening of the duration of a supergame. We addressed this question by conducting a supplementary experiment in which we reduced the duration to the very extreme: two round. Thus, in this experiment we consider the «lower bound setting» with the most unfavorable conditions for any kind of reputation building.

The experimental design and procedures of the supplementary experiment were identical to the main experiment, with the following exception: To keep the total number of rounds that subjects played in a session constant, each subject played six matches with two rounds each. After each match, the subjects were re-matched according to a round-robin (or «perfect stranger») scheme such that each pair of subjects played at most one match together. We omitted the «Baseline» (BSL) condition and implemented only a «Costless Monitoring» (CSM) and a «Costly Monitoring» (CYM) condition. Participants were recruited from the same general undergraduate student population of the University of Heidelberg. In total 60 subjects participated. Subjects were randomly assigned to treatment conditions, 36 in the CSM, 24 in the CYM condition.

First, we investigate whether and to what extent the key findings of section 4.2 also hold for two-round supergames. The analogue to figure 1 is shown in figure 7. Given the rematching scheme that we use in the supplementary experiment, the statistical test procedure needs to slightly adjusted: We cannot treat each match as an independent observation any more, since each subject played not only one but six matches in total. However, since roles remained constant, we can use the individual averages taken over all twelve rounds as an independent observation.[23]. Thus, all statistical tests reported in this section are Mann-Whitney rank-sum tests applied to

---

[23]A trustor never interacted with another trustor, and a trustee never interacted with another trustee. One may still object that the average behaviors of two trustors are in some kind of relationship mediated by behavior of their common trustees. However, this possible dependence appears sufficiently remote such that we dare to ignore it here, given that this experiment is only a supplement.

Figure 7: Results of the supplementary experiment by treatment condition. Panel (a) presents the aggregate results analogous to figure 1, panels (b) and (c) the dynamics in the CSM and the CYM conditions, respectively, analogous to figure 2, and panel (d) the frequency of monitoring as a fraction of all instances of trust in CYM, analogous to figure 3.

(a)

(b)

(c)

(d)

this time-averaged cross-section.

In the CSM treatment, trustors trusted 45.8 percent of the time (99 out of 216 cases) and in the CYM condition, 36.8 percent of the time (53 out of 144 cases). The difference is not statistically significant ($p = .4688$). Thus, the result that there is no less trust if trustworthiness is costly to monitor than under free observability still holds in two-round supergames.

In the CYM condition, the average trustor spent 7.2 tokens per match on monitoring, which is 19.8 percent of her gross payoff (36.4 tokens) and 43.0 tokens per session (over twelve periods). This is around twice as much as in the twelve-round supergames reported above. However, trustors' average payoffs (including monitoring expenditures) were not significantly lower (14.9 tokens) in the CYM condition than in the CSM condition (15.2 tokens, $p = .8987$). The same holds for trustees (23.0 vs. 26.2, $p = .1601$).[24]

This reinforces and sharpens our conclusion of section 4.2 that the treatment variations induced different monitoring behavior without affecting trust significantly. In addition, taken the results of the our main experiment and the supplementary experiment presented here together suggest that trustors tailor their monitoring expenditures to the structural features of the interaction in a way that they are approximately just as well off in a setting with monitoring costs than in a setting with free observability.

Comparing the CYM treatment with the CSM treatment in the supplementary experiment with respect to the information structure, we find a similar response to cost as in our main experiment, although the magnitudes differ somewhat. In the CYM condition trusting first movers monitored in about two thirds of the time (38 out of 53 cases, or 71.7 percent) and trusted without monitoring in the remaining cases. Despite the low number of observations (12 in the CYM condition, 18 in the CSM condition) , trust without monitoring is still marginally significantly more frequent in CYM than in CSM ($p = .0577$). Furthermore, comparing the CYM conditions between the main and the supplementary experiment, trustors trust without monitoring significantly more often in the former ($p = .0035$).[25]

To sum up, the key treatment effects obtained in our main experiment with twelve-round supergames qualitatively re-appear in the supplementary experiment with two-round supergames. Quantitatively, the longer supergames support on average more trust without monitoring than the shorter ones.

With respect to the match dynamics, figures 7b and 7c show the share of trustors that trusted with and without monitoring and the share of trustworthy trustees split up by round (analogous to figures 2b and 2c), and figure 7d the frequency of trust with and without monitoring, respectively, as a fraction of all instances of trust (analogous to figure 3). The intertemporal shift from trust with monitoring towards trust with-

---

[24]The difference is of non-negligible size and close to the statistical significance margin; it is possible that the difference gets statistically significant with a somewhat larger sample size.

[25]This Mann-Whitney test is applied to a dataset consisting of all trustors in the CYM condition from both experiments; each observation is the relative frequency of trust without monitoring over all twelve rounds of a session.

Figure 8: Dynamics of beliefs by treatment. Panel (a) depicts the average trustor's belief in their trustee acting trustworthy. Panel (b) depicts the average trustee's belief in being monitored.

**(a)**   **(b)**
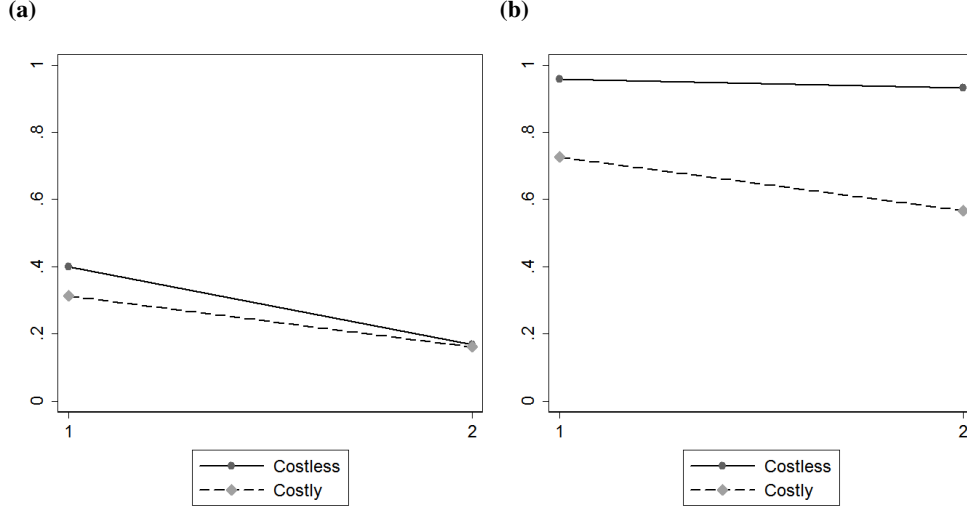
out monitoring is much less pronounced compared to the main experiment. There are two explanations for this. First, two rounds may be just too few for the kind of «model building» of the co-player explained above. Second, and perhaps more importantly, it may be an artifact of the experimental design. Given that trustors played six matches in a session of the supplementary experiment, information about their current trustee's terminal round action may still be valuable to them, because they could learn something about the population of trustees (despite they never interacting with this particular trustee again). Consistent with the latter explanation is the fact that only 8.3 percent of the trustors monitored in the 12th round of a session, while 30.0 percent monitored in terminal rounds of the initial five matches.

Analogous to figure 4, figure 8 summarizes information about the belief dynamics of trustors and trustees across rounds. The key properties of the patterns are consistent with the ones from the main experiment. Figure 8a shows the average belief of trustors in the trustworthiness of their co-player. As in the main experiment, they are very similar across both conditions. Also consistent with the main experiment, the average trustee in the supplementary experiment does anticipate that trustors monitor less, in particular in the second round, in the CYM condition than in the CSM condition. This is illustrated in figure 8b.

Figure 9 shows the realized payoffs over time in the two treatment conditions of the supplementary experiment. In the CSM condition trustors reaped on average on average 36.6 percent (36.0 percent in the first round, 37.3 percent in the second) of the joint payoff. As in the main experiment, in the CYM condition the average trustee reaped larger payoffs in the second half of the match (here round two) compared to

Figure 9: Average payoffs over time, in the CSM condition (solid) and the CYM condition (dashed).

the first half (here round one). In the first round, the average trustor received on average a share of 42.1 percent of the joint payoff (36.0 percent in the CSM condition), in the final round 36.0 percent (37.3 in the CYM condition).