

Bartling, Björn; Fehr, Ernst; Huffman, David; Netzer, Nick

Working Paper

The causal effect of trust

Working Paper, No. 304

Provided in Cooperation with:

Department of Economics, University of Zurich

Suggested Citation: Bartling, Björn; Fehr, Ernst; Huffman, David; Netzer, Nick (2018) : The causal effect of trust, Working Paper, No. 304, University of Zurich, Department of Economics, Zurich, <https://doi.org/10.5167/uzh-157382>

This Version is available at:

<https://hdl.handle.net/10419/192913>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



**University of
Zurich** ^{UZH}

University of Zurich
Department of Economics

Working Paper Series
ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 304

The Causal Effect of Trust

Björn Bartling, Ernst Fehr, David Huffman and Nick Netzer

October 2018

The Causal Effect of Trust

Björn Bartling, Ernst Fehr, David Huffman, and Nick Netzer*

October 17, 2018

Trust affects almost all human relationships – in families, organizations, markets and politics. However, identifying the conditions under which trust, defined as people’s beliefs in the trustworthiness of others, has a causal effect on the efficiency of human interactions has proven to be difficult. We show experimentally and theoretically that trust indeed has a causal effect. The duration of the effect depends, however, on whether initial trust variations are supported by multiple equilibria. We study a repeated principal-agent game with multiple equilibria and document empirically that an efficient equilibrium is selected if principals believe that agents are trustworthy, while players coordinate on an inefficient equilibrium if principals believe that agents are untrustworthy. Yet, if we change the institutional environment such that there is a unique equilibrium, initial variations in trust have short-run effects only. Moreover, if we weaken contract enforcement in the latter environment, exogenous variations in trust do not even have a short-run effect. The institutional environment thus appears to be key for whether trust has causal effects and whether the effects are transient or persistent.

Keywords: Trust, causality, equilibrium selection, belief distortions, incomplete contracts, screening, institutions

JEL: C91, D02, D91, E02

* Bartling, Fehr, Netzer: Department of Economics, University of Zurich, Blümlisalpstrasse 10, 8006 Zurich, Switzerland, e-mail: bjoern.bartling@econ.uzh.ch, ernst.fehr@econ.uzh.ch, nick.netzer@econ.uzh.ch. Huffman: Department of Economics, University of Pittsburgh, 4901 Wesley Posvar Hall, 230 South Bouquet Street, Pittsburgh, PA 15260, USA, e-mail: huffmand@pitt.edu.

1. Introduction

Almost every economic transaction involves incomplete contracts. Complexity and informational constraints render it impossible in many cases to govern all conceivable contingencies in a contract. In such circumstances, contracting parties may only be willing to interact and realize the associated efficiency gains, if they *trust* that the other party will not take advantage of them. It has therefore been argued that trust has fundamental implications for economic efficiency (see, e.g., Banfield, 1958; Arrow, 1972; Coleman, 1990; Putnam, 1993, 2000; Fukuyama, 1995). The potential scope for trust to shape economic outcomes is extremely broad. Trust could affect individual-level economic interactions, the efficiency of organizations, the functioning of entire markets, and even economic development and growth at the country level.

It is intuitive that differences in trust, which we define in this paper as *people's beliefs in the trustworthiness of others*, matter when contracts are incomplete. It is less straightforward, however, that these effects are *causal*. Differences in trust may well be caused by differences in economic conditions and outcomes, rather than causing these differences (see, e.g., the discussion in Algan and Cahuc, 2013 and 2014). This difficulty is exacerbated by the fact that it is hard to identify and measure all the subtle – formal and informal – institutional details that affect trust. Moreover, even if there is a causal effect of trust on economic outcomes, it is still not straightforward that this is more than a *short-run* effect. If there are opportunities for repeated interaction, and if learning about a contracting party's trustworthiness takes place, then trusting or distrusting beliefs should converge to their correct value over time. For initial trust differences to have *long-run* effects, it is necessary that these differences are self-confirming in some way.

The aim of this paper is (i) to empirically identify conditions under which trust does or does not have a causal impact on economic outcomes, (ii) to isolate when these effects are long-lasting and when they are only temporary, and (iii) to examine potential interactions between trust and the institutional environment. We address these questions by using data from laboratory experiments. A key advantage of using an experimental approach is the ability to induce exogenous variation in trust and to have full control over the institutional environment, which allows for a clean separation of the effects of trust and institutions and their possible interaction.

Our experiments involve principal-agent interactions, where principals make binding price offers and agents choose what quality to provide. Efficiency is increasing in quality, but there is a conflict of interest as higher quality benefits the principal while being costly for the agent.

Contracts are incomplete due to moral hazard in the agents' quality choice. There are opportunities to interact repeatedly over the course of 15 periods with different or the same trading partner, with noisy feedback about quality after each interaction. This experimental design is well suited for studying possible short- and long-run causal effects of trust because it features (i) incomplete contracts, such that there is a role for trust, and (ii) repeated play, such that we can observe how the efficiency of interactions evolves over time.

We implement variation in initial trust of principals through randomly assigning examples of real historical play that involved agents being either trustworthy or untrustworthy. The high-trust example involves agents responding to higher prices by choosing higher qualities. In the low-trust example, agents choose relatively low quality regardless of price. Providing examples of past play seems a natural approach, as one way that trust differences might arise is due to historical accident. A manipulation check confirms that principals updated their beliefs based on these examples, in the direction of expecting agents to choose high quality in response to high prices (high-trust) or to choose low quality regardless of price (low-trust).

We find that high and low levels of trust cause persistently high and low levels of economic efficiency in this environment, respectively. Principals in our treatment with the high-trust example pay high prices, agents choose high quality validating principals' trust (except for an end-game effect in the final periods), and efficiency is high. Principals in our treatment with the low-trust example, by contrast, pay low prices, agents choose low quality consistent with principals' beliefs, and efficiency remains low throughout the game.

To explain these results we turn to analyzing a theoretical model that captures key features of the experimental game. The model involves a mix of selfish and reciprocal agents. We show that there exists a high-trust sequential equilibrium that is a dynamic screening equilibrium: Principals pay high prices initially, continue to hire and pay high prices to an agent for whom the noisy feedback is positive, while paying low prices to an agent for whom the signal was negative. There also exists a low-trust sequential equilibrium that is a pooling equilibrium: Principals pay low prices, all agents choose low quality, so the noisy signal is uninformative, and principals do not learn anything that challenges their distrust.

We hypothesize that the positive long-run effect of initial trust arises in our experiment through the mechanism of equilibrium selection, i.e., by inducing the subjects to play the high-trust or the low-trust equilibrium. To test this explanation, we conducted additional treatments.

One set of treatments changes the strategic setting such that our model predicts that there should be a unique low-trust equilibrium. This involves eliminating the possibility for repeated interaction with the same trading partner, which precludes dynamic screening. We find that high trust still has a positive effect in such a setting, but the effect is only short-run. Principals begin by paying high prices, but agents are not trustworthy, and principals reduce prices over time.

We further hypothesize that this short-run effect reflects a learning process, in which principals only gradually realize that their beliefs were distorted and high prices are not rewarded by high quality in the unique equilibrium. To test whether the positive short-run effect of trust is in fact due to a belief distortion that is not supported by equilibrium behavior, we conducted a final set of treatments that again involve a unique low-trust equilibrium, but in which the principals' best response is independent of their beliefs about agents' types. Prices and quality are chosen simultaneously in these treatments, so that high prices cannot lead agents to choose high quality. In line with our hypothesis, principals in this final set of treatments pay low prices from the outset, regardless of whether they receive the high- or low-trust example.

Our paper complements a large empirical literature on trust and economic outcomes, which has mainly used survey data. Survey measures of trust have been shown to be positively correlated with a wide range of favorable economic outcomes, at the micro and macro level. For example, high trust is associated with individual-level willingness to invest, better functioning of credit markets, the ability of organizations to grow in size, international trade, and higher GDP per capita (e.g., Knack and Keefer, 1997; La Porta et al., 1997; Guiso et al., 2004, 2008a, 2009; Bloom et al. 2012). The causality underlying these correlations is, however, not entirely clear because good outcomes might lead to high trust. In addition, good outcomes and high trust might both be driven by some third factor, like strong institutions (e.g., North, 1981; Acemoglu et al., 2001; Acemoglu and Robinson 2012).

Efforts to address causality in the survey-based trust literature have primarily used instrumental variable (IV) approaches. In the seminal work of Knack and Keefer (1997), the authors use the extent of ethno-linguistic homogeneity, and the number of law students in 1963 as a percentage of all postsecondary students, as instruments for trust in regressions explaining growth rates in GDP. While it seems plausible that these variables are related to trust, it cannot be excluded that ethnic homogeneity and weak contract enforcement, indicated by many law students,

could directly affect economic growth, thereby invalidating the instruments.¹ Studies that are more recent have made progress by using alternative instruments. For example, Guiso et al. (2009) document a fascinating relationship between mutual trust between countries and the volume of trade, with common religion and somatic similarity as instruments for trust. A potential objection to the instruments, however, is that both could have a direct impact on trade.² Algan and Cahuc (2010) use lagged changes in trust over time to explain income per capita. This avoids a confound affecting many previous studies, that trust levels could stand in for time-invariant country characteristics like fixed institutions. There remains the question where changes in trust come from, and the possibility that they are caused by time-varying factors such as changes in the institutional environment.³ Our paper is complementary to this literature: While data from laboratory experiments can be challenged with regard to their external validity (but see, e.g., Herbst and Mas, 2015; Charness and Fehr, 2015; Kessler and Vesterlund, 2015), our approach has the important advantage that it allows implementing exogenous variation in trust and a tight control of the institutional environment. Our experimental approach thus provides a clean identification of the causal role of trust, the role of the institutional environment, and the possible interaction between the two.

There is also a related literature on trust using laboratory experiments. Previous papers in the gift-exchange literature (e.g., Fehr et al., 1993; Brandts and Charness, 2004; Brown et al. 2004; Charness, 2004; Charness et al., 2004) and more recent papers on the counterproductive effects of sanctions and other measures that constrain shirking by agents (e.g., Bohnet et al., 2001; Fehr and Rockenbach, 2003; Falk and Kosfeld, 2006; Bartling et al., 2012) suggest that trust might be self-confirming. However, this literature does not show that trust has a causal effect, because initial trust levels were endogenous. Costa-Gomes et al. (2014) is an exception in that they develop an

¹ Note that the use of lagged variables may not address the problems, because weak legal enforcement in 1963 is likely to be correlated with the state of legal enforcement far in the future.

² Under the *assumption* that one of the two instruments is valid (i.e., exogenous to the error term) a Hausman test does not reject the null hypothesis that both instruments are valid. This means we cannot reject the null hypothesis, but we also do not know whether it is true.

³ Specifically, the authors use changes in “inherited” trust levels of residents in the United States whose forbears immigrated to the country in different time-periods. While evidence exists that trust is transmitted across generations within families, this “inheritance” need not be genetic (the authors do not claim otherwise), and indeed transmission of trust has been shown to at least partly reflect socialization or imitation (Dohmen et al., 2012). Thus, changes in inherited trust over time may reflect changing social rather than genetic factors of parents, which could reflect changing institutions or other social trends in the country of origin. Other studies using instrumental variables include Guiso et al. (2006) and Tabellini (2010).

approach to instrumenting beliefs about the level of re-payment in a simultaneous-move version of the trust game proposed by Berg et al. (1995). Their paper provides the first evidence that exogenous belief variation can have a significant impact on choices. However, since they study a one-shot game, their paper does not address the question whether trust can become self-confirming, leading to sustainable long-run differences in economic performance.⁴

Finally, the existing theoretical literature has shown that different levels of trust can arise in a given economic environment due to multiple equilibria (e.g., Tabellini, 2008; Aghion et al., 2010) or multiple stable long-run outcomes of dynamic learning processes (e.g., Bower et al., 1996; Guiso et al., 2008b; Aghion et al., 2011). Our experiment provides a first clean test of this general idea. Our theoretical model, however, differs from the literature in two important ways. First, we follow a standard game-theoretic approach with fixed preferences, while Tabellini (2008) and Aghion et al. (2010) study behavior that is transmitted from generation to generation and coevolves slowly with external institutions. Our theoretical and empirical results show that trust is malleable rather quickly and can have immediate causal effects that persist in the long-run, even with fixed preferences and institutions. Second, in models like Bower et al. (1996), where agents learn about a given population state, the long-run levels of trust and economic efficiency cannot be manipulated by interventions that select between different equilibria. By contrast, we show that selecting the right equilibrium is an important consideration in the design of organizations and mechanisms.

The remainder of the paper is organized as follows. In Section 2, we explain our experimental design and present a manipulation check showing that our exogenous variation of trust is effective. Section 3 demonstrates that trust can have a causal long-run impact on economic outcomes. Section 4 contains a theoretical analysis of the principal-agent game, suggesting that the long-run impact of trust may be due to equilibrium selection. In Section 5, we present evidence for the proposed mechanism of equilibrium selection, and we show that trust can still have a causal short-run effect even if the equilibrium is unique. Section 6 concludes with a discussion of the strong complementarity between trust and institutions in achieving good economic outcomes.

⁴ Also related, Bohnet and Huck (2004) test whether old institutions “have an afterglow.” They study a repeated one-shot sequential trust game and show that experience with reputation institutions have lingering positive effects on trustworthiness of second-movers, but not on trust of first-movers, after reputation possibilities are removed. Trustworthiness (and trust) is declining steadily over time, however, so the effect seems to dissipate in the long-run.

2. Experimental Design

Our main research question is whether exogenously manipulated beliefs of principals about the agents' trustworthiness can have a *causal* impact on economic outcomes. Moreover, we are interested in the *duration* of the effect. We therefore adopt a principal-agent framework in which contracts are incomplete, so that trust may matter for the efficiency of the interaction. Moreover, we allow for repeated interaction, so that we can study how efficiency evolves over time.

A typical principal-agent relation is plagued with a conflict of interest. While a higher quality level by the agent increases the principal's expected value, providing high quality is costly for the agent. Our specific design builds on the experimental markets in Brown et al. (2004) but extends their framework in an important way. In our design, principals can only observe an informative signal about the agents' quality choices and not their quality choices directly, as in Brown et al. (2004). This feature adds realism, in that it is not possible in many types of economic interactions to identify the role of agents' effort versus luck in determining outcomes. The signal is observable by the principal and the agent but it is not verifiable by third parties and therefore not contractible, giving rise to moral hazard. The principal's belief that an agent is trustworthy may then be relevant for determining the efficiency of the market interaction. We define agents to be trustworthy when their equilibrium behavior is responsive to prices. Trustworthy agents reciprocate a high price offer with a high quality choice in equilibrium. Untrustworthy agents, in contrast, always choose low quality levels in equilibrium, irrespective of the offered prices.

One important feature of many markets is the ability to repeatedly interact with the same contracting party. Long-term trading relations are possible in our main treatments because subjects have fixed identification numbers over the course of the experiment and the principals' contract offers can be addressed to specific agents. Therefore, a principal can make contract offers to the same agent in consecutive periods and, if the agent accepts the offers, a long-term relation can be established. We will argue in Section 4 that a long-run impact of exogenous variation in the principals' initial trust then materializes because different levels of trust correspond to different equilibria of the repeated market game. In one equilibrium, all agents choose low quality levels irrespective of the offered prices. In the other equilibrium, only selfish agents choose low quality levels, while reciprocal agents respond in a trustworthy way and choose high quality when price offers are high. We will also show that exogenous variations in trust can have a positive effect in

settings with a unique equilibrium, but this effect arises only in the short-run because the principals' high levels of trust are not supported by the agents' equilibrium behavior.

2.1 Market Game

The principals are the contract makers, i.e., they alone can make contract offers to agents, who can choose among the available offers; agents themselves cannot make offers to principals. There are 15 trading periods. Each period, a principal can purchase at most one unit of the traded good from an agent, and an agent can accept at most one contract offer. A period has two stages. In stage one, contract offers are made. If a principal and an agent conclude a contract, they enter stage two, where the agent has to determine a quality level.

A contract offer consists of a price offer $p \in \{1, \dots, 100\}$, a desired quality level $\tilde{q} \in \{1, \dots, 9\}$, and the principal's identification number (ID). The price offer p is third-party verifiable and thus contractible, while the lack of third-party verifiability of actual quality rules out that quality levels are contractually enforceable. There are two different types of offers, public and private. In private offers, a principal also indicates the agent's ID with whom he wants to trade, and then only this agent is informed about the contract offer. The possibility to make private contract offers enables principals to engage in long-term relationships. In public offers, all agents (and also the other principals) are informed about the offer; hence, each agent has the chance to accept a public offer. A principal can make as many private offers and as many public offers as he wants in a given period. However, once an agent accepts one of the offers, the principal is matched with this agent, learns the ID of the matched agent, and his other outstanding offers are removed from the market. To prevent principals from making private offers to agents who have already concluded a contract with another principal, principals are at all times informed about which agents remain in the market. The default at the beginning of each period is that no agent has a contract and no principal has made an offer. There are always ten agents and seven principals in a market, i.e., there is an excess supply of three agents.

The principal has to pay the offered price p if an agent accepts a contract, but the agent can choose any actual quality level $q \in \{1, \dots, 9\}$, irrespective of the desired quality level $\tilde{q}$, because actual quality is only observable by the agent but not by the principal, and thus not contractible. The agent's quality choice stochastically determines the value the principal receives. For simplicity, there are only two possible value levels in our experiments, 100 and 10. The probability

that the principal receives the high value is given by $(q/10)$, while with probability $1 - (q/10)$ the principal receives the low value. Thus, if an agent chooses the lowest possible quality level of 1, then the probability that the principal receives the high value is 10 percent; if the agent chooses the highest possible quality level of 9, the probability is 90 percent. Hence, the actual quality q can never be inferred perfectly from the realization of the value.

The (expected) material payoffs of principals and agents in a period are given by

$$E[\Pi^{principal}] = \begin{cases} 100 \cdot \frac{q}{10} + 10 \cdot \left(1 - \frac{q}{10}\right) - p & \text{if a contract is concluded} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

$$\Pi^{agent} = \begin{cases} p - c(q) & \text{if a contract is concluded} \\ 5 & \text{otherwise} \end{cases} \quad (2)$$

where $c(q)$ denotes the cost of providing quality. The outside option of an agent who does not conclude a contract is 5. Table 1 shows the cost function $c(q)$ of the agents. The cost function is strictly increasing and exhibits weakly increasing marginal costs. Since the marginal cost of quality is at most 3, while the marginal expected revenue is always 9, the efficient quality level is given by $q = 9$. The payoff functions (1) and (2), the number of principals and agents, the cost function $c(q)$, and the fact that there are 15 trading periods is common knowledge.

Table 1: Agents' Cost Function

Quality	1	2	3	4	5	6	7	8	9
Cost	0	1	2	4	6	8	10	12	15

At the end of each period, each subject is informed about the contract $(p, \tilde{q})$ he had concluded, his own payoff, as well as about the trading partner's ID. An agent is also informed about his current principal's payoff. A principal, however, is not informed about his current agent's payoff, because a principal does not observe the agent's quality choice and thus the cost of providing this quality level. The subjects write this information on a printed form that is provided along with the experimental instructions. This procedure ensures that each subject can always remind herself about her own trading history.

2.2 Inducing Variation in Principals' Trust

To exogenously vary the principals' trust levels, we randomly assigned them to two different information conditions. In the high-trust treatment, the principals were informed about a "historical example" in which agents behaved in a very trustworthy manner; in the low-trust treatment, they were shown an example in which the agents displayed a low level of trustworthiness. More specifically, for our high-trust treatments we selected the market from Brown et al. (2004) that had the steepest price-quality relation, and for the low-trust treatments we selected the market with the flattest price-quality relation.

The example was provided at the end of the experimental instructions. Subjects were informed that the information provided was an "example," and that it showed how quality is related to price "in a past session." Subjects were told that the information in the example was something that they "could use in their decisions today." The description of the source of the example was accurate but deliberately vague, and we did not claim that the information provided about a single past session was representative.

Figure 1 shows how we presented the examples to the subjects in the instructions.⁵ The top row was shown to the principals in the high-trust treatments, the bottom row to the principals in the low-trust treatments. On the left, the price-quality relation is shown. The figure shows the average quality provided by the agents in the example for each of the given bins of offered prices. On the right, we show how this price-quality relation translates into a price-payoff relation, given the principals' payoff function in our experiment. The high-trust example involved agents being trustworthy, in that they strongly responded to high prices with high quality levels. In the low-trust example, agents were untrustworthy; they provided rather low quality for all price levels.


Note that the examples contain no information about the historical frequency of price choices by principals. They do indicate the range of prices that was used, but this was identical across the high-trust and low-trust examples. This is deliberate, to rule out that the examples influence behavior by conveying information about historical behavior of principals. Rather, the differential information content across examples is solely about the trustworthiness of agents. Any impact should thus come through the beliefs of principals about trustworthiness. We will examine

⁵ We provide the experimental instructions in the Appendix.

below the extent to which our trust manipulation was effective in the sense that it differentially affected principals' beliefs.

Subjects in the role of agents did not receive any example, nor were they informed that the subjects in the role of principals received such information. The instructions for agents thus did not differ in the high- and in the low-trust treatments, which rules out any direct impact on outcomes through an influence on agents. This illustrates the advantages of an experimental setting for varying only the principals' trust, defined in our context as the principals' belief in the trustworthiness of agents.

Figure 1: The High-Trust and Low-Trust Examples Shown to Principals.



Notes: The top row shows the example provided to the principals in the high-trust treatments, the bottom row shows the example provided in the low-trust treatments. The price-quality relation is shown on the left, the corresponding expected price-payoff relation on the right.

2.3 Treatment Conditions and Experimental Procedures

We conducted five markets for each of our two main treatments, denoted HT and LT, where HT and LT stand for high- and low-trust, respectively. We also conducted markets for various control treatments. These include two markets of a treatment denoted HT-Long, where we implement the same conditions as in HT, except that the game lasted 25 periods rather than 15 periods. We use HT-Long to clarify the possible role of end-game effects. Moreover, we conducted treatments HT-R and LT-R, with five markets each, where R stands for random IDs. These involved the same conditions as HT and LT, except that subject IDs were randomly reshuffled every period, so that repeated interaction between the same parties were ruled out and only one-shot interactions were possible. In the theoretical part of our paper, we will argue that multiple equilibria no longer exist in the absence of repeated interactions. Thus, the purpose of the HT-R and LT-R treatments is to study the causal role of trust in a unique equilibrium environment. Another pair of control treatments, HT-R2 and LT-R2, had random IDs and two-sided contractual incompleteness, where R stands for random IDs and 2 stands for two-sided contractual incompleteness. Not only were the agents free to choose any actual quality level, but also the principals did not have to pay the offered price p ; actual quality levels and prices were chosen simultaneously in these treatments. In the theoretical part of our paper, we predict that exogenous changes in trust should not even have short-run causal effects when choices are simultaneous. The HT-R2 and LT-R2 treatments, therefore, serve the purpose of testing this prediction. Table 2 provides an overview of our treatments.

Table 2: Treatment Overview

Treatment	# Markets	Trust Example	Identification Numbers	Contractual Incompleteness	# Periods
HT	5	high	fixed	one-sided	15
LT	5	low	fixed	one-sided	15
HT-R	5	high	random	one-sided	15
LT-R	5	low	random	one-sided	15
HT-R2	5	high	random	two-sided	15
LT-R2	5	low	random	two-sided	15
HT-Long	2	high	fixed	one-sided	25

We implemented a between-subjects design, i.e., each subject participated in only one market in one treatment. Altogether we have 32 markets, with seven principals (buyers) and ten agents (sellers) each. Hence, 544 subjects participated in our experiment. Subjects were mainly students from the University of Zurich and the Swiss Federal Institute of Technology in Zurich. Students majoring in economics or psychology were not eligible to participate. Since the terms “principal” and “agent” are not in common usage among student subjects, the experiment was framed in terms of “buyers” and “sellers.”

All sessions took place at the computer laboratory of the Department of Economics at the University of Zurich. The study was computerized with the software z-Tree (Fischbacher, 2007) and the recruitment was conducted with the software ORSEE (Greiner, 2015). Before the subjects entered the lab, they randomly drew a place card that specified at which computer terminal to sit. The terminal number determined a subject’s role as either principal or agent, which remained fixed throughout the experiment.

Subjects received written instructions including comprehension questions, which had to be answered correctly before a session could begin. A summary of the instructions was read aloud by the experimenter with the aim of generating common knowledge of the instructions. There were also two practice periods before the actual experiment to make the subjects familiar with the market procedures. Subjects only went through the first stage of the experiment in both practice periods, so that principals did not observe payoffs and could not draw inferences about agents’ actual quality choices. No money could be earned during the two practice periods.

Sessions lasted about 2.5 hours. Payoffs from the experiment, denominated in points, were converted into money at the rate of 10 points to CHF 1 (about \$ 1.05 at the time of the experiments) at the end of a session. On average, subjects earned about CHF 47.65, which includes a show-up fee of CHF 20. The subjects received their payments privately.


2.4 Manipulation Check

Our experimental approach aims at exogenously varying the principals’ beliefs about the agents’ trustworthiness. Figure 2 provides a manipulation check by showing the principals’ expectations about the empirical relationship between offered prices and agent quality. These expectations were elicited at the beginning of the experiment, after reading the instructions but before entering the

trading periods. We asked principals to predict what they thought would be the average quality level chosen by agents, conditional on different possible offered prices.

The figure reveals that our belief manipulation was successful in all three treatment pairs, HT versus LT, HT-R versus LT-R, and HT-R2 versus LT-R2. Principals expected significantly higher average quality levels across the range of prices when they had received the high-trust example rather than the low-trust example, and also expected significantly steeper relationships between price and quality. Regressions confirm that the differences in average expected quality were statistically significant at the 1-percent level, as were the differences in slopes, in all three treatment pairs.⁶ We cannot reject the hypotheses that the average expected quality, and the slopes of the price-quality relations, are identical when comparing across treatments involving the low-trust example, and across treatments involving the high-trust example.⁷

Figure 2: Manipulation Check



Notes: The black lines show the price-quality relation that principals expect in the three high-trust treatments (HT, HT-R, and HT-R2). The grey lines show expectations in the three low-trust treatments (LT, LT-R, and LT-R2).

⁶ The results for average expected quality are from OLS regressions of principals' expectations about quality on the relevant treatment dummy, clustering standard errors at the subject level. All treatment dummy coefficients are significant at the 1-percent level. Results for the differences in slopes are from OLS regressions of expected quality on the relevant treatment dummy, the offered price, and an interaction term, clustering on subject. All coefficients of the interaction terms are significant at the 1-percent level.

⁷ The results are based on OLS regressions of principals' expectations about quality on the appropriate treatment dummies and interaction terms, clustering standard errors at the subject level.

Since the agents do not receive historical examples about agents' trustworthiness in a previous experiment, successful random assignment should result in no differences in agents' beliefs across treatments involving high-trust versus low-trust examples. Indeed, we cannot reject the hypothesis that the agents' expectations about average quality, and the slope of the price-quality relations, are identical within each of the treatment pairs.⁸ Since the agents indicate their "homegrown" beliefs, we can compare these beliefs with the beliefs that the principals indicate. We find that the principals' beliefs in the high-trust treatments roughly correspond to the agents' homegrown beliefs. Principals who received the low-trust example have beliefs that are more pessimistic than homegrown beliefs.

3. Long-Run Causal Effects of Trust

In this section, we present the results of our main treatments HT and LT and show that the differences in historical information about agents' trustworthiness lead to significant long-run differences in market outcomes. In Section 4, we will provide a game-theoretic analysis of our market game, in which the possibility to screen reciprocal from selfish agents generates multiple equilibria, and historical information affects outcomes through equilibrium selection. Our additional treatment pairs HT-R versus LT-R and HT-R2 versus LT-R2 will then allow us to test additional predictions of the theory and to verify the mechanism for a causal long-run impact of trust.

3.1 Trusting Behavior of Principals

The manipulation check in Section 2.4 confirmed that random assignment to HT versus LT causally leads to a significant difference in initial beliefs of principals about the trustworthiness of agents. We now turn to analyzing the impact of these beliefs on prices paid by principals. All else equal, higher prices indicate more trusting behavior by principals, since quality is not contractually enforceable.


Figure 3 shows the average price levels over time in HT and LT. The figure reveals that initial high beliefs in trustworthiness in HT are associated with a persistent level of high trusting behavior, whereas the lower initial beliefs in trustworthiness in LT are associated with a persistent

⁸ The results are based on OLS regressions of agents' expectations about quality on the appropriate treatment dummies and interaction terms, clustering standard errors at the subject level.

low level of trusting behavior. Average price offers were 43.32 in HT and 17.69 in LT, a difference that is highly significant (Wilcoxon rank sum test, $p < 0.01$). Regression analysis indicates that this treatment difference is stable over time; there is no statistically significant time trend for the treatment difference in prices (see regressions (1) and (4) in Table A1 in the Appendix). We summarize these observations in our first result.

Result 1: *The two different historical examples of agents' trustworthiness induce significant and stable differences in principals' price setting behavior. Favorable information about agents' trustworthiness leads principals to pay prices that are more than twice as high on average in HT than in LT, where they receive unfavorable information about agents' trustworthiness.*

Figure 3: Price Levels Over Time in Treatments HT and LT




Notes: The black lines show average prices in sessions with the high-trust example and the grey lines show average prices in sessions with the low-trust example.

3.2 Agents' Quality Choices and Market Efficiency

Agents' quality choices are a crucial outcome variable as they determine market efficiency in our set-up. The large and stable difference in prices in HT and LT, respectively, suggests that agents may be reacting to these different price levels in a way that creates a feedback that reinforces the principals' initial beliefs.

Figure 4 shows the average quality levels over time in HT and LT. The figure reveals that quality is in fact substantially higher in HT than in LT, in line with the treatment difference in price offers. Average quality was 5.46 in HT and 3.32 in LT, a difference that is highly significant (Wilcoxon rank sum test, $p < 0.01$). The difference appears stable over the course of the game, except for an end-game effect that emerges in the last couple of periods. A regression analysis confirms that there is no statistically significant time trend for the treatment difference in qualities, excluding the final two periods (see regression (4) in Table A2 in the Appendix). Thus, principals in HT, expecting high quality and offering high prices, ended up receiving high quality and observing a high frequency of project success. Principals in LT, expecting low quality and offering low prices, ended up receiving low quality and observing a low frequency of project success.

Figure 4: Quality Levels Over Time in Treatments HT and LT



Notes: The black lines show average quality in sessions with the high-trust example and the grey lines show average qualities in sessions with the low-trust example.

To check whether the decline in quality towards the end of the game was only an end-game effect, and did not instead reflect a systematically fading effect of the principals' initial high belief in agents' trustworthiness, we conducted two markets of control treatment HT-Long. Treatment HT-Long is identical to treatment HT, except that it lasted 25 periods rather than 15 periods. Figure A1 in the Appendix shows that the decline in quality was indeed an end-game effect. Quality remained high for 10 additional periods in HT-Long, compared to HT, and started to decline only as the game approached its end. Increasing the length of the game thus simply moved the decline to the last periods of the longer game.⁹ We summarize these observations in our next result.

Result 2: *Agents respond to the principals' higher level of trusting behavior in HT relative to LT with significantly higher quality levels in HT.*

In sum, high and low initial beliefs about agents' trustworthiness lead to sustained levels of high and low trusting behavior by principals, respectively. High and low levels of trusting behavior, in turn, bring about high and low quality choices by agents, which reinforce the initial beliefs about trustworthiness. Thus, there appear to be high-trust and low-trust equilibria, characterized by high and low market efficiency, which are selected by the historical examples of agents' trustworthiness. In the next section, we provide a game-theoretic analysis, which supports our interpretation that two equilibria coexist in our principal-agent market game.

4. Theory and Mechanism

In this section, we summarize our theoretical analysis of the principal-agent market game. The formal analysis can be found in the Appendix. To keep the analysis tractable, the game that we solve is a simplified version of the game used in the experiment. An essential feature of the experimental setting is the finite repetition of the stage game, coupled with an excess supply of agents. This allows principals to rehire or fire agents conditional on the stochastic outcome of their earlier interaction. We capture the dynamic interaction with just two periods. The feature of excess supply of agents is modelled by having one principal and two agents. Furthermore, the game is simplified by assuming that all actions (prices and qualities) are binary.

⁹ Treatment HT-Long replicates the high level of efficiency observed in HT. Average quality in HT-Long is 5.57, which is almost identical to (and not significantly different from) the average quality of 5.46 in HT.

While the principal is profit-maximizing, we assume that there are two types of agents, selfish and reciprocal.¹⁰ Only reciprocal types find it optimal to respond to a high price with high quality already in a one-shot interaction. However, since the agents' types are not observable to the principal, and since we assume that the share of reciprocal types is not too large, offering the high price is not profitable for the principal in a one-shot interaction. The one-shot game thus has a unique equilibrium in which price and quality are always low.

With repeated interaction, by contrast, the game exhibits coexistence of a low-trust and a high-trust sequential equilibrium for a large range of parameters, including parameters that closely resemble the payoff structure in the experiment.

The low-trust equilibrium replicates the outcome of the one-shot interaction. The principal initially offers the low price to one of the agents, and both types of that agent respond with low quality. The subsequent stochastic realization of the value is therefore not informative about the agent's type. As a consequence, the principal again offers the low price in the second period, and both types of the agent respond with low quality. A reciprocal agent cannot signal his type by a first-period deviation from equilibrium because the quality choice is not directly observable to the principal. Thus, the low-trust equilibrium is a pooling equilibrium in which gift-exchange between the principal and the reciprocal types does not materialize.

In the high-trust equilibrium, the principal initially trusts an agent of unknown type, that is, he pays the high price. A selfish type responds with low quality but a reciprocal type responds with high quality. The principal's belief that the agent is a reciprocal type declines if he receives the low value. He will then not offer the high price again in the second period. A realized high value, by contrast, constitutes a positive signal about the type of the agent. Given the positively updated belief, the principal's expected profit is maximized by offering the high price again to the same agent. A selfish type has no incentive to mimic a reciprocal type because he does not obtain the additional intrinsic benefit from responding to high price with high quality. Thus, the high-trust equilibrium is a separating equilibrium, where the initial trusting behavior of the principal serves to stochastically screen reciprocal types from selfish types.

¹⁰ One could assume that the principal is reciprocal, too, but this matters less because the principal is the first mover. Reciprocity may provide an additional reason for the principal to pay a high price in the second period only if the stochastic signal indicates that the agent is a reciprocal type. However, we show below that this is also an equilibrium strategy for a selfish principal.

We interpret the historical information about agents' trustworthiness used in the experiment as a device that selects between these multiple equilibria. The predictions of the high-trust and the low-trust equilibrium are indeed in line with the experimental findings in HT and LT, respectively.

Consider the low-trust equilibrium first. The offered price and the returned quality are predicted to be low in both periods on the equilibrium path, as confirmed by the experimental results in LT. Furthermore, if the principal trembled and offered the high price in the first period, then both types of the agent would still respond with low quality in that equilibrium. This is sustained by the correct off-equilibrium belief that the principal subsequently reacts to the realized value (which is still uninformative because both types behave in the same way) by not rehiring the agent. The agent's off-equilibrium behavior thus confirms the principal's belief about a flat price-quality reaction as induced by the historical example in LT.

Consider the high-trust equilibrium next. The high first-period price elicits an average quality strictly above the low level, as confirmed by our results in HT. If the principal trembled and offered the low price instead, both types of the agent would respond with low quality in this equilibrium. This confirms the principal's belief in a positive price-quality reaction as induced by the historical example in HT. A response of quality to price can also be observed on the equilibrium path of the high-trust equilibrium. Depending on the stochastic realization of the value, the principal offers either the high or the low price in the second period, and the induced expected quality is larger in the former case than in the latter.

5. Short-Run Causal Effects of Trust

Our explanation for the causal long-run impact of trust on economic outcomes is a mechanism of equilibrium selection. We therefore hypothesize that a long-run impact of trust should not exist in an environment with a unique equilibrium. In this section, we examine this hypothesis and show empirically that the effect of inducing high trust among the principals is indeed not stable and declines significantly over time when the equilibrium is unique. However, there is a short-run effect of trust that is driven by temporary belief distortions among the principals. These distorted beliefs are generated by the initial high-trust manipulation that is not supported by equilibrium behavior.

5.1 Control Treatment 1 - Testing for Equilibrium Selection

To provide a direct test of the equilibrium selection hypothesis, we conducted additional treatments, denoted HT-R and LT-R. These treatments are identical to HT and LT, except that subjects' IDs are randomly reassigned every period rather than remaining fixed. Random reassignment of IDs rules out long-term relationships and implies that interactions are one-shot. In the context of our model, this change implies that there is a unique equilibrium, in which price and quality are always low.

Intuitively, the high trust equilibrium does not exist when interactions are one-shot, because the trusting behavior of the principal in the high-trust equilibrium serves to separate reciprocal types from selfish types. Separation enables the principal to offer the high price again to the same agent only when the realization of the value indicates that the agent is a reciprocal type with sufficiently high probability. However, screening is a costly investment for the principal in the model, because the high first-period price is not reciprocated often enough to cover the cost instantaneously. Its value lies in the information that it generates about the agent's type, and this value can only be monetized if the principal is able to make a targeted offer to that agent again in the future. Thus, if equilibrium selection is the mechanism underlying the impact of trust, then the prediction is that the historical examples about agents' trustworthiness have no lasting effect in the one-shot environment and outcomes in HT-R should be similar to outcomes in LT-R.

Panel (b) of Figure 5 shows the price levels over time in HT-R and LT-R. In contrast to the environments with repeated interactions (shown again in Panel (a) to allow for easy comparisons across all treatments), the historical examples about agents' trustworthiness are associated with a weaker and less stable difference in prices. Prices are 33.88 on average in HT-R and 16.45 in LT-R. This difference of about 17 points is still statistically significant (Wilcoxon rank sum test, $p < 0.01$) but substantially smaller than the difference of 26 points observed in treatments HT and LT with fixed IDs. The average price offers in LT and LT-R are not significantly different (17.69 vs. 16.45, Wilcoxon rank sum test, $p = 0.81$) and they remain similar throughout the whole 15 periods. In contrast, the price pattern in HT and HT-R is very different. While prices in HT-R are initially only slightly lower than in HT, the gap between these treatments strongly increases over time because prices in HT-R are steadily declining. Regression analysis shows that the overall impact of the historical examples on prices is significantly weaker in treatment pair HT-R versus LT-R than in treatment pair HT versus LT (see the interaction term between the dummy variables

for the high-trust example and for treatment pair HT-R/LT-R in regression (1) of Table A3 in the Appendix; $p < 0.02$). Furthermore, the price difference between HT-R and LT-R becomes much smaller over time, showing that the initial impact of the historical examples is steadily declining in a setting without repeated interactions (the interaction term between a dummy variable for the high-trust example and period is negative and statistically significant in regressions (2) and (5) in Table A1 in the Appendix). We summarize these observations in our next result.


Result 3: *The two different historical examples of agents' trustworthiness induce sizeable initial differences in principals' price setting behavior in HT-R and LT-R, where interactions are one-shot, but these price differences steadily decline over time.*

We now turn to the behavior of agents. Recall that in the presence of multiple equilibria there are stable quality differences over time (shown again in Panel (a) of Figure 6), except for an end-game effect. In contrast, Panel (b) of Figure 6 shows that the difference in quality levels between treatments HT-R and LT-R is declining over time. This time trend is significant as indicated by regression analysis and is driven by the decline in quality levels in HT-R (see regressions (2) and (5) in Table A2 of the Appendix). The steady decline in quality in HT-R, in combination with relatively flat quality levels over time in LT-R, has the consequence that the difference in the average quality between HT-R and LT-R (1.40 points) is smaller than the difference between HT and LT (2.10 points).¹¹ Regression analysis confirms that the impact of the historical high-trust example is significantly weaker in treatment pair HT-R versus LT-R than in treatment pair HT versus LT (see the interaction term between the dummy variables for the high-trust example and for treatment pair HT-R vs. LT-R in regression (3) of Table A4, which excludes the end-game effect; $p < 0.08$). The next result summarizes this pattern.

Result 4: *Agents respond to principals' initially higher levels of trusting behavior in HT-R relative to LT-R with higher quality levels in HT-R. However, differences in quality levels between HT-R and LT-R are smaller than between HT and LT and declining over time.*


¹¹ Average quality is 3.68 in HT-R and 2.28 in LT-R (Wilcoxon rank-sum test, $p = 0.03$).

Figure 5: Price Levels Over Time in All Treatments



Notes: The black lines show average prices in sessions with the high-trust example and the grey lines show average prices in sessions with the low-trust example.

Figure 6: Quality Levels Over Time in All Treatments



Notes: The black lines show average qualities in sessions with the high-trust example and the grey lines show average qualities in sessions with the low-trust example.

In sum, the favorable historical example about agents' trustworthiness in HT-R still leads to higher market efficiency compared to LT-R, although our theoretical analysis predicts a unique low-trust equilibrium. However, the effect fades over time, suggesting that the observed difference may be a temporary distortion away from equilibrium rather than an equilibrium phenomenon. In particular, it seems plausible that the favorable historical example distorts upwards some principals' beliefs about the prevalence of reciprocal types. We should then expect these principals to offer high prices initially. It takes some time for them to observe that the high price is not reciprocated frequently enough to be profitable in the one-shot setting, where the benefits from investment in screening for agent type cannot be realized. This should generate learning dynamics towards the actual equilibrium, which appears consistent with the observed data. We will investigate this hypothesis more rigorously in the next subsection.

5.2 Control Treatment 2 - Testing for Belief Distortions

To investigate whether belief distortions may have caused the short-run effect of the high-trust and low-trust examples in the treatments with random IDs, we implemented a third treatment pair, denoted HT-R2 and LT-R2. In addition to ruling out repeated interaction, we also weaken the enforceability of contracts by making the principals' price offers non-binding. Specifically, although the principal stipulates a price offer and a desired quality level like in the other treatments, the price offer is non-binding and the principal is free to pay any actual price $p \in \{1, \dots, 100\}$ irrespective of the initially offered price. Hence, the contractual incompleteness is two-sided: Not only are the agents' quality levels not third-party enforceable, also the principals' price offers are not enforced by third parties. Once an agent has accepted a contract with a non-binding price offer and a desired quality level, principals and agents choose the actual price and the actual quality simultaneously.

The simultaneity of choices makes it impossible to induce higher quality by paying higher prices. Our theoretical analysis therefore implies that the principal's best response is to pay a low price even if the favorable historical example about agents' trustworthiness induces a distorted belief about the strength of reciprocity, because the principal has no incentive to act on this belief. Therefore, we predict that in HT-R2 prices will be significantly lower than in HT-R and, as a consequence, the price difference between the high- and the low-trust condition will be much smaller in treatment pair HT-R2 versus LT-R2 than in treatment pair HT-R versus LT-R.

Furthermore, the lower price difference between HT-R2 and LT-R2 will be associated with lower quality differences compared to treatment pair HT-R versus LT-R.

Panel (c) of Figure 5 shows actual prices paid over time in HT-R2 and LT-R2. A comparison with panel (b) shows that average prices in HT-R2 are much lower than in HT-R, suggesting that the high-trust manipulation indeed had a much lower effect in treatment HT-R2. Moreover, although there is initially a small price difference between HT-R2 and LT-R2 this difference quickly vanishes completely.¹² This contrasts with the price difference in treatment pair HT-R versus LT-R, where the belief distortion led to a prolonged and much larger price difference. Regression analysis confirms that the impact of the trust examples is significantly weaker in treatment pair HT-R2 versus LT-R2 than in treatment pair HT-R versus LT-R (see regressions (2) and (4) in Table A3 in the Appendix; $p < 0.01$). Furthermore, the difference in prices between HT-R2 and LT-R2 is getting significantly smaller over time (see regressions (3) and (6) of Table A1 in the Appendix). We summarize these observations in our next result.

Result 5: *The two different historical examples of agents' trustworthiness induce only very small and vanishing differences in principals' price setting behavior in HT-R2 and LT-R2, where interactions are one-shot and contractual incompleteness is two-sided. This contrasts with the larger behavioral differences in treatment pair HT-R versus LT-R, where principals had an incentive to act on their belief distortions.*

We turn next to the behavior of the agents. Panel (c) of Figure 6 shows that the high-trust condition has no effect on agents' behavior as their quality levels are basically identical in treatments HT-R2 and LT-R2 right from the beginning, such that there are also no average quality differences (2.51 and 2.59, Wilcoxon rank-sum test, $p = 0.75$). This pattern differs from treatment pair HT-R versus LT-R, where the principals' belief distortion appears to have led to (slowly declining) price differences, which then translated into (slowly declining) quality differences between the high- and the low-trust condition. Regression analyses confirm that the difference in quality levels across high- and low-trust conditions in treatment pair HT-R2 versus LT-R2 is highly significantly

¹² Figure A2 in the Appendix shows offered and actual prices in treatments HT-R2 and LT-R2. Offered prices start out high and actually increase over time, but actual price payments are much lower and decline over time. Actual average prices are 11.61 in HT-R2 and 8.47 in LT-R2 (Wilcoxon rank-sum test, $p = 0.08$).

different from those in treatment pair HT-R versus LT-R (see regressions (2) and (4) of Table A4 in the Appendix; $p < 0.01$). Our final result summarizes these observations.

Result 6: *The two different historical examples of agents' trustworthiness have no effect on quality levels in HT-R2 relative to LT-R2. Quality levels in HT-R2 and LT-R2 are indistinguishable and converge toward the minimal level.*

In sum, we hypothesized that principals' belief distortions explain the short-run effect of the historical examples in treatments HT-R and LT-R. To test this hypothesis we developed the treatment pair HT-R2 versus LT-R2, in which principals' beliefs about agents' trustworthiness should not affect their price offers. The data confirm this prediction as the historical examples have basically no effects in treatments HT-R2 and LT-R2, whereas they induced (slowly declining) price and quality differences in treatment pair HT-R versus LT-R.

6. Conclusions

It is well-documented that high levels of trust are correlated with favorable economic outcomes, such as growth, trading volume, gains from trade, the functioning of organizations, or overall welfare of social groups (e.g., Knack and Keefer, 1997; La Porta et al., 1997; Guiso et al., 2004, 2008a, 2009; Bloom et al. 2012). Establishing whether trust has a *causal* effect on such outcomes, however, has been challenging (e.g., Fehr, 2009; Algan and Cahuc 2014). Moreover, the mechanisms by which trust can affect economic outcomes are unclear. In this paper, we provided evidence from laboratory experiments that trust can affect economic outcomes in a causal and sustained way. We generated an exogenous instrument for trust by varying historical examples of past play that involved agents either being trustworthy or untrustworthy in a repeated gift-exchange game (treatments HT and LT). Our manipulation check confirmed that the historical examples influenced principals' beliefs about agents' trustworthiness. The market game data showed that the historical examples significantly affected principals' trusting behavior. A stable equilibrium with high prices and high product quality emerged when principals were provided with good information about agents' trustworthiness. In contrast, principals and agents were trapped in an inefficient equilibrium with low prices and low quality when principals were provided with bad information about agents' trustworthiness. Our data thus show that—in the same economic

environment—high levels of trust can generate good economic outcomes, while low levels of trust can generate bad economic outcomes.

We also designed additional treatment pairs with random IDs (HT-R vs. LT-R) and with two-sided contractual incompleteness (HT-R2 vs. LT-R2). In both treatment pairs, a unique bad equilibrium exists such that equilibrium selection cannot occur. We indeed found that our initial trust manipulation does not cause stable long-run differences in outcomes but triggers only short-run effects (HT-R vs. LT-R) or no effects at all (HT-R2 vs. LT-R2), consistent with the proposed mechanisms of equilibrium selection for long-run effects and temporary belief distortions for short-run effects.

The additional treatment pairs also reveal, more generally, that the impact of trust depends on the economic environment in which people interact. Trust only operates in the long-run in favorable environments, characterized by (partial) contract enforcement and the ability to enter into long-term relationships (as in our main treatments HT and LT). In less favorable environments, where repeated interaction is not possible (as in HT-R and LT-R) and, on top, contract enforcement is weaker (as in HT-R2 and LT-R2), a unique unfavorable equilibrium exists and a positive long-run impact of trust does not arise. Hence, our data suggest that good and stable economic outcomes depend on having the right combination of high trust and favorable environmental factors.

Formal institutions like the rule of law are a major factor that shapes the economic environment, for instance by determining whether contracts can be enforced (e.g., La Porta et al., 2008). Formal institutions can also affect the economic environment by influencing the extent to which long-term relationships are sustainable, for instance by creating political and economic stability and holding individuals accountable for their past actions (e.g., Acemoglu and Robinson, 2012). Our empirical results therefore suggest a strong complementarity between trust and formal institutions. At least two implications follow from this argument.

First, policies aimed at improving economic performance could remain ineffective if they focus only on the quality of formal institutions. As our results show, even with good institutions low trust can lead to low levels of economic performance. On the other hand, measures that aim at moving society out of a low-trust trap—such as public awareness campaigns that promote codes of conduct or advertise role models of trustful business relations—may remain ineffective if formal

institutions are weak. Our results suggest that reforms are more likely to be successful if they improve formal institutions *and* take care of selecting the right equilibrium at the same time.

Second, future research on the determinants of economic performance should focus more on interactions between separate factors of influences. Controlling for trust and institutional quality, but not for their interaction, might yield results that obfuscate the real effects. Similar issues may arise, more generally, for the interaction between environmental factors and informal norms of behavior. This raises a large range of novel questions for future research.

References

- Acemoglu, Daron, Simon Johnson, and James A. Robinson (2001). "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, 91(5), 1369-1401.
- Acemoglu, Daron and James A. Robinson (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown Publishing Group.
- Aghion, Philippe, Yann Algan, Pierre Cahuc, and Andrei Shleifer (2010). Regulation and Distrust. *Quarterly Journal of Economics*, 125(3), 1015-1049.
- Aghion, Philippe, Yann Algan, and Pierre Cahuc (2011). "Civil Society and the State: The Interplay Between Cooperation and Minimum Wage Regulation," *Journal of the European Economic Association*, 9(1), 3-42.
- Algan, Yann and Pierre Cahuc (2010). "Inherited Trust and Growth," *American Economic Review*, 100(5), 2060-2092.
- Algan, Yann and Pierre Cahuc (2013). "Trust and Growth." *Annual Review of Economics*, 5, 521-549.
- Algan, Yann and Pierre Cahuc (2014). "Trust, Growth, and Well-Being: New Evidence and Policy Implications." In: *Handbook of Economic Growth*, Vol. 2A, Chapter 2, 1st Edition, eds. Philippe Aghion and Steven Durlauf, p. 49-120.
- Arrow, Kenneth J. (1972). "Gifts and Exchanges," *Philosophy and Public Affairs*, 1(4), 343-362.
- Banfield, Edward C. (1958). *The Moral Basis of a Backward Society*. Free Press.
- Bartling, Björn, Ernst Fehr, and Klaus M. Schmidt (2012). "Screening, Competition, and Job Design: Economic Origins of Good Jobs," *American Economic Review*, 102(2), 834-864.
- Berg, Joyce, John Dickhaut, and Kevin McCabe (1995). "Trust, Reciprocity, and Social History," *Games and Economic Behavior*, 10, 122-142.
- Bloom, Nicholas, Raffaella Sadun, John Van Reenen (2012). "The Organization of Firms across Countries," *Quarterly Journal of Economics*, 127(4), 1663-1705.
- Bohnet, Iris, Bruno S. Frey, and Steffen Huck (2001). "More Order with Less Law: On Contract Enforcement, Trust, and Crowding," *American Political Science Review*, 95, 131-144.
- Bohnet, Iris and Steffen Huck (2004). "Repetition and Reputation: Implications for Trust and Trustworthiness When Institutions Change," *American Economic Review*, 94(2), 362-366.

- Bower, Anthony B., Steven Garber, and Joel C. Watson (1996). "Learning about a population of agents and the evolution of trust and cooperation," *International Journal of Industrial Organization*, 15, 165-190.
- Brandts, Jordi, and Gary Charness (2004). "Do Labour Market Conditions Affect Gift Exchange? Some Experimental Evidence," *Economic Journal*, 114, 684-708
- Brown, Martin, Armin Falk, and Ernst Fehr (2004). "Relational Contracts and the Nature of Market Interactions," *Econometrica*, 72, 747-780.
- Charness, Gary (2004). "Attribution and Reciprocity in an Experimental Labor Market," *Journal of Labor Economics*, 22(3), 665-88.
- Charness, Gary and Ernst Fehr (2015). "From the lab to the real world," *Science*, 350, 512-513.
- Charness, Gary, Guillaume R. Frechette, and John H. Kagel (2004). "How Robust Is Laboratory Gift Exchange?" *Experimental Economics*, 7(2), 189-205.
- Coleman, James (1990). *Foundations of Social Theory*. Harvard University Press.
- Costa-Gomes, Miguel A., Steffen Huck, and Georg Weizsäcker (2014). "Beliefs and actions in the trust game: Creating instrumental variables to estimate the causal effect," *Games and Economic Behavior*, 88, 298-309.
- Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2012). "The intergenerational transmission of risk and trust attitudes," *The Review of Economic Studies*, 79(2), 645-677.
- Falk, Armin, and Michael Kosfeld (2006). "The Hidden Costs of Control," *American Economic Review*, 96, 1611-1630.
- Fehr, Ernst (2009). "On the Economics and Biology of Trust." *Journal of the European Economic Association*, 7(2-3), 235-266.
- Fehr, Ernst, Georg Kirchsteiger, and Arno Riedl (1993). "Does Fairness Prevent Market Clearing—An Experimental Investigation," *Quarterly Journal of Economics*, 108, 437-459.
- Fehr, Ernst, and Bettina Rockenbach (2003). "Detrimental effects of sanctions on human altruism," *Nature*, 422, 137-140.
- Fischbacher, Urs (2007). "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics*, 10(2), 171–78.
- Fukuyama, Francis (1995). *Trust: The Social Virtues and the Creation of Prosperity*. Free Press, New York.

- Greiner, Ben (2015). "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 1(1), 114-125.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2004). "The Role of Social Capital in Financial Development," *American Economic Review*, 94, 526-556.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2006). "Does Culture Affect Economic Outcomes?" *Journal of Economic Perspectives*, 20, 23-48.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2008a). "Trusting the Stock Market," *Journal of Finance*, 63(6), 2557-2600.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2008b). "Social Capital as Good Culture," *Journal of the European Economic Association*, 6(2-3), 295-320.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2009). "Cultural Biases in Economic Exchange?" *Quarterly Journal of Economics*, 124(3), 1095-1131.
- Herbst, Daniel and Alexandre Mas (2015). "Peer effects on worker output in the laboratory generalize to the field," *Science*, Vol. 350, 545-549.
- Kessler, Judd, and Lise Vesterlund (2015). "The external validity of laboratory experiments: the misleading emphasis on quantitative effects." In: *Handbook of Experimental Economic Methodology*, eds. Guillaume Fréchette and Andrew Schotter, p. 391-406.
- Knack, Stephen, and Philip Keefer (1997). "Does Social Capital Have an Economic Payoff? A Cross-Country Investigation," *Quarterly Journal of Economics*, 112, 1251-1288.
- La Porta, Rafael, Florencio López-de-Silanes, Andrei Shleifer, and Robert Vishny (1997). "Trust in Large Organizations." *American Economic Review*, 87(2), 333-38.
- La Porta, Rafael, Florencio López-de-Silanes, and Andrei Shleifer (2008). "The Economic Consequences of Legal Origins." *Journal of Economic Literature*, 46(2), 285-332.
- North, Douglass C. (1981). *Structure and Change in Economic History*. Norton.
- Putnam, Robert D. (1993). *Making Democracy Work*. Princeton UP, Princeton, NJ.
- Putnam, Robert D. (2000). *Bowling Alone: The Collapse and Revival of American Community*. Simon and Schuster, New York.
- Tabellini, Guido (2008). "The Scope of Cooperation: Values and Incentives," *Quarterly Journal of Economics*, 123(3), 905-950.
- Tabellini, Guido (2010). "Culture and Institutions: Economic Development in the Regions of Europe," *Journal of the European Economic Association*, 8(4), 677-716.

Appendix

A. Additional Tables and Figures

Table A1: Time Trends in Prices Offered by Principals as a Function of Trust Example and Treatment

	All periods			Periods < 14		
	HT vs. LT	HT-R vs. LT-R	HT-R2 vs. LT-R2	HT vs. LT	HT-R vs. LT-R	HT-R2 vs. LT-R2
	(1)	(2)	(3)	(4)	(5)	(6)
High-trust	26.83*** [3.61]	23.06*** [2.33]	8.09*** [2.08]	25.76*** [3.59]	23.49*** [2.52]	7.88*** [2.36]
Period	-0.01 [0.19]	-0.02 [0.15]	-0.64*** [0.13]	-0.02 [0.21]	-0.02 [0.19]	-0.65*** [0.15]
High-trust * Period	-0.14 [0.25]	-0.71** [0.28]	-0.61*** [0.22]	0.08 [0.28]	-0.79** [0.33]	-0.57** [0.29]
Constant	17.69*** [2.87]	16.66*** [1.52]	13.58*** [1.33]	17.73*** [2.82]	16.66*** [1.72]	13.60*** [1.40]
Observations	1042	1046	1031	904	906	898

Notes: Panel regression estimates, with random effects for principals, clustering standard errors on principal. Columns (1) to (3) present regressions for the three different treatment pairs, respectively, using all periods. Columns (4) to (6) do the same, except for excluding the final two periods. Columns (1) and (4) use only data from treatment pair HT vs. LT treatments, columns (2) and (5) use only data from treatment pair HT-R vs. LT-R, and columns (3) and (6) use only data from treatment pair HT-R2 vs. LT-R2. The omitted category is the respective low-trust treatment. “High-trust” is a dummy variable indicating the respective high-trust treatment. “Period” takes on values 1 to 15, indicating the respective period. ***, **, * denote significance at the 1, 5, and 10 percent level, respectively.

Table A2: Time Trends in Quality Levels Chosen by Agents as a Function of Trust Example and Treatment

	All periods			Periods < 14		
	HT vs. LT	HT-R vs. LT-R	HT-R2 vs. LT-R2	HT vs. LT	HT-R vs. LT-R	HT-R2 vs. LT-R2
	(1)	(2)	(3)	(4)	(5)	(6)
High-trust	3.10*** [0.50]	2.06*** [0.35]	0.76* [0.39]	2.78*** [0.50]	2.09*** [0.38]	0.66 [0.42]
Period	-0.03 [0.03]	-0.00 [0.02]	-0.11*** [0.03]	0.02 [0.03]	-0.01 [0.02]	-0.13*** [0.03]
High-trust * Period	-0.12*** [0.04]	-0.08** [0.03]	-0.09** [0.04]	-0.06 [0.05]	-0.09** [0.04]	-0.07 [0.04]
Constant	3.46*** [0.39]	2.29*** [0.22]	3.40*** [0.28]	3.20*** [0.36]	2.32*** [0.22]	3.651*** [0.30]
Observations	1042	1046	1031	904	906	898

Notes: Panel regression estimates, with random effects for principals, clustering standard errors on principal. Columns (1) to (3) present regressions for the three different treatment pairs, respectively, using all periods. Columns (4) to (6) do the same, except for excluding the final two periods. Columns (1) and (4) use only data from treatment pair HT vs. LT treatments, columns (2) and (5) use only data from treatment pair HT-R vs. LT-R, and columns (3) and (6) use only data from treatment pair HT-R2 vs. LT-R2. The omitted category is the respective low-trust treatment. “High-trust” is a dummy variable indicating the respective high-trust treatment. “Period” takes on values 1 to 15, indicating the respective period. ***, **, * denote significance at the 1, 5, and 10 percent level, respectively.

Table A3: Impact of Historical Examples on Prices, Difference-in-Differences

	All periods		Periods < 14	
	(1)	(2)	(3)	(4)
High-trust	25.73*** [3.32]	17.42*** [1.39]	26.28*** [3.33]	17.94*** [1.37]
R-Treatments	-1.1 [2.68]		-1.04 [2.70]	
High-trust * R-Treatments	-8.34** [3.60]		-8.37** [3.60]	
R2-Treatments		-8.01*** [0.90]		-7.42*** [0.93]
High-trust * R2-Treatments		-14.27*** [1.60]		-14.13*** [1.64]
Constant	17.58*** [2.56]	16.46*** [0.77]	17.57*** [2.59]	16.50*** [0.76]
Observations	2088	2077	1810	1804


Notes: Panel regression estimates, with random effects for principals, clustering standard errors on principals. The sample is restricted to treatments HT, LT, HT-R and LT-R for columns (1) and (3), and to HT-R, LT-R, HT-R2, and LT-R2 for columns (2) and (4). Column (1) gives the differential impact of varying the historical example in treatment pair HT-R vs. LT-R and treatment pair HT vs. LT. The omitted category in columns (1) and (3) is treatment LT. Column (2) gives the differential impact of varying the historical example in treatment pair HT-R2 vs. LT-R2 and in treatment pair HT-R vs. LT-R. The omitted category in columns (2) and (4) is treatment LT-R. “High-trust” is a dummy variable indicating a high-trust treatment. “R-Treatments” is a dummy variable indicating treatments HT-R and LT-R. “R2-Treatments” is a dummy variable indicating treatments HT-R2 and LT-R2. Columns (3) and (4) do the same as columns (1) and (2) but exclude the final periods. ***, **, * denote significance at the 1, 5, and 10 percent level, respectively.

Table A4: Impact of Historical Examples on Quality Levels, Difference-in-Differences

	All periods		Periods < 14	
	(1)	(2)	(3)	(4)
High-trust	2.13*** [0.43]	1.40*** [0.19]	2.35*** [0.44]	1.48*** [0.21]
R-Treatments	-0.94*** [0.36]		-1.06*** [0.38]	
High-trust * R-Treatments	-0.73 [0.47]		-0.87* [0.49]	
R2-Treatments		0.23 [0.16]		0.30 [0.18]
High-trust * R2-Treatments		-1.34*** [0.26]		-1.30*** [0.28]
Constant	3.22*** [0.35]	2.28*** [0.11]	3.35*** [0.36]	2.28*** [0.12]
Observations	2088	2077	1810	1804


Notes: Panel regression estimates, with random effects for principals, clustering standard errors on principals. The sample is restricted to treatments HT, LT, HT-R and LT-R for columns (1) and (3), and to HT-R, LT-R, HT-R2, and LT-R2 for columns (2) and (4). Column (1) gives the differential impact of varying the historical example in treatment pair HT-R vs. LT-R and treatment pair HT vs. LT. The omitted category in columns (1) and (3) is treatment LT. Column (2) gives the differential impact of varying the historical example in treatment pair HT-R2 vs. LT-R2 and in treatment pair HT-R vs. LT-R. The omitted category in columns (2) and (4) is treatment LT-R. “High-trust” is a dummy variable indicating a high-trust treatment. “R-Treatments” is a dummy variable indicating treatments HT-R and LT-R. “R2-Treatments” is a dummy variable indicating treatments HT-R2 and LT-R2. Columns (3) and (4) do the same as columns (1) and (2) but exclude the final periods. ***, **, * denote significance at the 1, 5, and 10 percent level, respectively.

Figure A1: Robustness Check Verifying End-Game Effect in Final Periods



Notes: Average quality levels are more volatile over the course of the experiment in HT-Long than in HT because we conducted only two markets in HT-Long, not five as in HT.

Figure A2: Offered Prices and Actual Prices with Two-Sided Contractual Incompleteness



B. Game-Theoretic Analysis

The game described below is simpler than the game played in the experiment, but captures its essential features. We model the dynamic interaction by considering two periods and the excess supply of agents by considering one principal and two agents. The firm is assumed to be profit-maximizing, while agents can be either selfish types or reciprocal types. Finally, we simplify the strategy space by assuming that all actions are binary.

We will be interested in environments in which the share of reciprocal types is not large enough to generate equilibrium gift-exchange in a one-shot interaction. We then show that multiple equilibria can exist in the dynamic interaction, one of them with gift-exchange and the other without. We interpret the historical information about agents' trustworthiness used in the experiment as a device that selects between these equilibria.

The Principal-Agent Problem

A principal offers a price $p \in \{p_L, p_H\}$ to an agent, where $0 \leq p_L < p_H$. After observing the price, the agent responds by producing a good of quality $q \in \{q_L, q_H\}$, where $0 < q_L < q_H < 1$. The good is either valuable, in which case it generates a payoff of v for the principal, or it is useless and does not generate any value. Quality q is the probability that the good is valuable. Denote by $\omega \in \{0, 1\}$ the state of the world describing whether the good is valuable ($\omega = 1$) or not ($\omega = 0$). The agent's cost of providing low quality is normalized to zero; the cost of providing high quality is c . Given actions (p, q) , the expected material payoffs of principal and agent are, respectively,

$$\pi^P(p, q) = qv - p \text{ and } \pi^A(p, q) = p - \left(\frac{q - q_L}{q_H - q_L} \right) c.$$

We assume that $0 < c < (q_H - q_L)v$, which implies that providing the high quality is efficient.

The principal is profit-oriented and maximizes $u^P(p, q) = \pi^P(p, q)$. The agent has a type $\alpha \in \{0, a\}$, where $0 < a < 1$, and maximizes

$$u^A(p, q, \alpha) = \pi^A(p, q) - \alpha |\pi^P(p, q) - \pi^A(p, q)|.$$

Type $\alpha = 0$ is selfish and cares only about own material payoff. Type $\alpha = a$ is inequity-averse, where the symmetric formulation of inequity-aversion is the simplest way of modelling a reciprocal motive.¹ The prior probability of the agent being reciprocal is given by $0 < \lambda < 1$.

¹ Note that this agent dislikes inequality in expected payoffs. One could also model aversion to the expectation of inequality in ex-post payoffs, after the stochastic value of the good has realized.

Let $q^*(p, \alpha) \in \arg\max_q u^A(p, q, \alpha)$ denote an optimal quality choice of the α -type agent in response to price p . For the selfish type, we obtain $q^*(p_L, 0) = q^*(p_H, 0) = q_L$. Trustworthy behavior of the reciprocal type arises if $q^*(p_L, a) = q_L$ and $q^*(p_H, a) = q_H$. The following assumption makes sure that this is indeed the case.

Assumption 1 (Trustworthiness)

$$u^A(p_L, q_L, a) > u^A(p_L, q_H, a) \text{ and } u^A(p_H, q_H, a) > u^A(p_H, q_L, a).$$

Given Assumption 1, let $p^*(\beta) \in \arg\max_p u^P(p, \beta q^*(p, a) + (1 - \beta)q^*(p, 0))$ denote an optimal offer of a principal who expects the high price to be reciprocated with probability $\beta \in [0, 1]$. We will be interested in environments where $p^*(\lambda) = p_L$, so that gift-exchange does not arise in a one-shot game when the principal's belief is given by the prior $\beta = \lambda$. However, gift-exchange becomes possible in a dynamic game, where the principal might be able to update her belief about the agent. Suppose there was an initial stage at which the selfish agent chooses q_L , while the inequity-averse agent chooses q_H . Then, if the good turns out to be of high value, a simple application of Bayes' rule implies that the principal's posterior belief would increase to

$$\bar{\beta} = \lambda \left(\frac{q_H}{\lambda q_H + (1 - \lambda)q_L} \right).$$

In addition to $p^*(\lambda) = p_L$, we will assume that $p^*(\bar{\beta}) = p_H$. Written in terms of the primitive parameters, this can be summarized as follows.

Assumption 2 (Value of Information)

$$\lambda v < \frac{p_H - p_L}{q_H - q_L} < \lambda \left(\frac{q_H}{\lambda q_H + (1 - \lambda)q_L} \right) v.$$

One-Shot Game

The timing of the one-shot interaction is given as follows:

1. The principal chooses price p .
2. Nature determines the agent's type α ($\alpha = a$ with independent probability λ).
3. The agent chooses quality q .
4. Nature determines the state ω ($\omega = 1$ with independent probability q).

The terminal nodes of this game are given by $t = (p, \alpha, q, \omega)$. The players' payoffs in the terminal nodes are

$$U^P(t) = u^P(p, \omega) \text{ and } U^A(t) = u^A(p, q, \alpha).$$

We assume that the agent observes the price p and her own type α , so the game has complete information. The principal's strategy prescribes the price to be chosen in the root of the game. The agent's strategy prescribes a quality to be chosen for each price-type combination (p, α) .

The one-shot game corresponds to treatments HT-R and LT-R in the experiment, where agents' random IDs prevent long-run relations. We obtain the following immediate result, the proof of which is left to the reader.

Proposition 1. *Under Assumptions 1 and 2, the one-shot game has a unique sequential equilibrium. In this equilibrium, the principal pays the low price and both types of the agent respond with low quality.*

The treatment pair HT-R2 and LT-R2 corresponds to a variant of the one-shot game where principal and agent make their choices simultaneously. It is straightforward to see that this variant of the game also has a unique equilibrium in which price and effort are low. This already holds under Assumption 1 alone, as paying the low price is a dominant strategy for the principal when choices are simultaneous. The fact that the principal's belief about the agent's type matters with sequential choices but not with simultaneous choices constitutes a difference between the treatment pair HT-R and LT-R on the one hand and the treatment pair HT-R2 and LT-R2 on the other hand. Suppose the high-trust example initially distorts upwards the beliefs of some principals about the share λ of trustworthy types. In treatment HT-R, we should then expect some principals to offer high prices initially. It takes some time for them to learn that the high price is not reciprocated frequently enough to be profitable, generating a slow learning dynamics towards the actual equilibrium. In treatment HT-R2, by contrast, distorted beliefs would not translate into high price offers because making a low price offer is the dominant strategy in this game. Hence, we should expect the actual equilibrium to be reached quicker with two-sided than with one-sided contractual incompleteness.

Dynamic Game

The timing of the dynamic interaction is given as follows:

1. The principal chooses price p^1 .
2. Nature determines agent 1's type α_1 ($\alpha_1 = a$ with independent probability λ).
3. Agent 1 chooses quality q^1 .
4. Nature determines the state ω^1 ($\omega^1 = 1$ with independent probability q^1).
5. The principal chooses whether to keep agent 1 ($k = 1$) or to fire her and hire agent 2 instead ($k = 0$). The principal also chooses price p^2 for the second period.
6. Nature determines agent 2's type α_2 ($\alpha_2 = a$ with independent probability λ).
7. The hired agent chooses quality q^2 .
8. Nature determines the state ω^2 ($\omega^2 = 1$ with independent probability q^2).

For notational simplicity, we assume that nature determines agent 2's type α_2 even if the agent is not hired in the second period. The terminal nodes of the dynamic game are then given by $t = (p^1, \alpha_1, q^1, \omega^1, k, p^2, \alpha_2, q^2, \omega^2)$. The players' payoffs are

$$U^P(t) = u^P(p^1, \omega^1) + u^P(p^2, \omega^2),$$

$$U_1^A(t) = u^A(p^1, q^1, \alpha_1) + ku^A(p^2, q^2, \alpha_1),$$

$$U_2^A(t) = (1 - k)u^A(p^2, q^2, \alpha_2).$$

Note that we assume here that the inequity-averse agents compare themselves only with the principal, separately period by period, whenever they interact.

Concerning the information structure, we assume that an agent's type and quality choice is observable only to the agent herself, while everything else is observable to all players.² A (pure) strategy of the principal prescribes the price to be chosen in the root of the game, $s^P(\emptyset) \in \{p_L, p_H\}$, as well as for each observed history (p^1, ω^1) a hiring decision and the price offered in the second period, $s^P(p^1, \omega^1) \in \{0, 1\} \times \{p_L, p_H\}$. A strategy of agent 1 prescribes a quality to be chosen in the first period for each observed price-type combination, $s_1^A(p^1, \alpha_1) \in \{q_L, q_H\}$, and a quality to be chosen conditional on all observables in case she is hired again in the second period,

² We could also assume that price offers are only observable to the currently hired agent, and/or that the realized value of the good is observable only to the principal. This would complicate the notation of beliefs, but we would still obtain the equilibrium outcomes derived below.

$s_1^A(p^1, \alpha_1, q^1, \omega^1, 1, p^2) \in \{q_L, q_H\}$. Finally, a strategy of agent 2 prescribes a quality to be chosen conditional on all observables in case she is hired in the second period, $s_2^A(p^1, \omega^1, 0, p^2, \alpha_2) \in \{q_L, q_H\}$. For each of the observable histories at which a player acts, she maintains a probabilistic belief over the nodes in the corresponding information set, i.e., a belief about the earlier unobservable actions that led to this information set.

The dynamic game corresponds to treatments HT and LT in the experiment, where agents' fixed IDs allow for repeated interaction. In a first step, we describe conditions under which the game admits a sequential equilibrium that replicates the one-shot outcome: The principal initially pays the low price and both types of the first agent respond with low quality; the principal then always fires the first agent and offers the low price to the second agent, who responds with low quality. We refer to such an equilibrium as a *low-trust equilibrium*. It can exist if the principal correctly believes that a high price would not elicit high quality from any type of the first agent, and hence would also not facilitate learning about that agent's type. It will turn out that the binding constraint for this construction is the reciprocal first agent's incentive not to respond to a high price with high quality. The following assumption makes sure that this constraint can be satisfied.

Assumption 3 (Low-Trust Incentive-Compatibility)

$$(q_H - q_L)u^A(p_L, q_L, a) > u^A(p_H, q_H, a) - u^A(p_H, q_L, a).$$

We can now state the following result.

Proposition 2. *Under Assumptions 1, 2, and 3 the dynamic game has a low-trust sequential equilibrium.*

Proof: We first characterize agent 2's strategy in any sequential equilibrium. After observing any history $(p^1, \omega^1, 0, p^2, \alpha_2)$ she entertains a belief about (α_1, q^1) , which must be consistent with the requirements imposed by sequential equilibrium. However, her optimal behavior does not depend on these beliefs. Under Assumption 1, we always obtain the unique sequentially rational choice $s_2^A(p^1, \omega^1, 0, p^2, \alpha_2) = q^*(p^2, \alpha_2)$. We next characterize agent 1's second-period strategy in any sequential equilibrium. After observing any history $(p^1, \alpha_1, q^1, \omega^1, 1, p^2)$ she entertains a belief about α_2 , which must be consistent with the requirements imposed by sequential equilibrium.

However, her optimal behavior does not depend on these beliefs. Under Assumption 1, we obtain the unique sequentially rational choice $s_1^A(p^1, \alpha_1, q^1, \omega^1, 1, p^2) = q^*(p^2, \alpha_1)$.

We now subsume these choices directly into the players' payoff functions and treat the game as a reduced game between the principal and agent 1. It ends in the terminal nodes $\hat{t} = (p^1, \alpha_1, q^1, \omega^1, k, p^2)$ with payoffs

$$\begin{aligned} U^P(\hat{t}) &= u^P(p^1, \omega^1) + ku^P(p^2, q^*(p^2, \alpha_1)) \\ &\quad + (1 - k)[\lambda u^P(p^2, q^*(p^2, a)) + (1 - \lambda)u^P(p^2, q^*(p^2, 0))], \\ U_1^A(\hat{t}) &= u^A(p^1, q^1, \alpha_1) + ku^A(p^2, q^*(p^2, \alpha_1), \alpha_1). \end{aligned}$$

This reduced game has two proper subgames, one starting after each possible first period price offer. In each of these subgames, the only non-singleton information sets are those of the principal when observing (p^1, ω^1) , where she entertains beliefs about (α_1, q^1) . Since $q_L > 0$ and $q_H < 1$, these beliefs can always be determined by Bayes' rule when we start from the root of the respective subgame. This uniquely pins down the consistent beliefs in any sequential equilibrium.

Consider first the subgame starting after $p^1 = p_L$. Let the strategies in this subgame be given by

$$s_1^A(p_L, 0) = s_1^A(p_L, a) = q_L \text{ and } s^P(p_L, 0) = s^P(p_L, 1) = (0, p_L),$$

i.e., both types of the first agent respond with low quality, and, irrespective of the realized value of the good, the principal then hires the second agent and pays the low price. Given any observation of (p_L, ω^1) , the principal entertains a probabilistic belief about (α_1, q^1) , but only the marginal distribution of α_1 matters for her sequentially rational choices (since q^1 is not payoff relevant conditional on ω^1 , and later behavior also does not depend on q^1). Denoting the probability attached to $\alpha_1 = a$ by $\beta^P(p_L, \omega^1)$, we obtain $\beta^P(p_L, 0) = \beta^P(p_L, 1) = \lambda$ from Bayes' rule. It then follows immediately from Assumptions 1 and 2 that the principal's strategy is indeed sequentially rational. As for the agent, observe that deviations cannot affect the principal's second period behavior. It then follows from Assumption 1 that the agent's strategy is also sequentially rational. The resulting expected payoff of the principal in the root of this subgame is $U_L^P = 2u^P(p_L, q_L)$.

Consider now the subgame starting after $p^1 = p_H$. Let the strategies in this subgame be given by

$$s_1^A(p_H, 0) = s_1^A(p_H, a) = q_L \text{ and } s^P(p_H, 0) = (1, p_L), s^P(p_H, 1) = (0, p_L),$$

i.e., both types of the first agent respond with low quality and the principal always pays the low price in the second period, keeping the first agent if and only if the good is of low value. We obtain the beliefs $\beta^P(p_H, 0) = \beta^P(p_H, 1) = \lambda$. Under Assumptions 1 and 2, the principal thus wants to pay the low price in the second period and is indifferent between keeping and firing the agent, which makes her strategy sequentially rational. As for the agent, consider type $\alpha_1 = 0$ first. Assumption 1 implies that q_L maximizes her first-period payoff. Moreover, $u^A(p_L, q^*(p_L, 0), 0) = p_L \geq 0$ implies that the selfish agent (weakly) benefits from a larger probability of being hired again in the second period, which implies that her strategy is sequentially rational. Consider next type $\alpha_1 = a$, who faces a trade-off between her payoff-maximizing response in the first period and the probability of being hired again in the second period. The condition for $s_1^A(p_H, a) = q_L$ to be sequentially rational is

$$\begin{aligned} u^A(p_H, q_L, a) + (1 - q_L)u^A(p_L, q^*(p_L, a), a) \\ \geq u^A(p_H, q_H, a) + (1 - q_H)u^A(p_L, q^*(p_L, a), a), \end{aligned}$$

which is satisfied under Assumptions 1 and 3. The resulting expected payoff of the principal in the root of this subgame is $U_H^P = u^P(p_H, q_L) + u^P(p_L, q_L)$.

Given the strategies and payoffs in the two subgames, it follows that $s^P(\emptyset) = p_L$ is the sequentially rational first-period price for the principal. ■

Next, we describe conditions under which the game admits an equilibrium in which gift-exchange occurs. The principal initially pays the high price, to which a selfish agent responds with low quality and a reciprocal agent responds with high quality. The principal then always keeps the agent but offers the high price in the second period if and only if the good turns out to be valuable. Thus, the principal tries to screen the reciprocal types from the selfish types. We refer to such an equilibrium as a *high-trust equilibrium*. Several constraints have to be satisfied for this equilibrium to exist, which we summarize in the following.

Assumption 4 (High-Trust Incentive-Compatibility)

- (i) $u^A(p_H, q_L, 0) - u^A(p_H, q_H, 0) > (q_H - q_L)[u^A(p_H, q_L, 0) - u^A(p_L, q_L, 0)],$
- (ii) $u^A(p_H, q_H, a) - u^A(p_H, q_L, a) > (q_H - q_L)[u^A(p_L, q_L, a) - u^A(p_H, q_H, a)],$
- (iii) $u^P(p_H, \lambda q_H + (1 - \lambda)q_L) + \lambda q_H u^P(p_H, q_H) + (1 - \lambda)q_L u^P(p_H, q_L) \\ > [2 - \lambda(1 - q_H) - (1 - \lambda)(1 - q_L)] u^P(p_L, q_L).$

We can now state the following result.

Proposition 3: *Under Assumptions 1, 2, and 4, the dynamic game has a high-trust sequential equilibrium.*

Proof: Consider again the reduced game between the principal and agent 1 constructed in the proof of Proposition 2. Also, let the strategies and beliefs in the subgame starting after $p^1 = p_L$ be the same as in the proof of Proposition 2, i.e.,

$$s_1^A(p_L, 0) = s_1^A(p_L, a) = q_L \text{ and } s^P(p_L, 0) = s^P(p_L, 1) = (0, p_L),$$

where $\beta^P(p_L, 0) = \beta^P(p_L, 1) = \lambda$, with a resulting expected payoff for the principal in the root of this subgame of $U_L^P = 2u^P(p_L, q_L)$.

Consider now the subgame starting after $p^1 = p_H$. Let the strategies be given by

$$s_1^A(p_H, 0) = q_L, s_1^A(p_H, a) = q_H \text{ and } s^P(p_H, 0) = (1, p_L), s^P(p_H, 1) = (1, p_H),$$

i.e., the selfish agent responds with low quality and the trustworthy agent responds with high quality, while the principal always keeps the agent but pays the high price in the second period only if the good turns out to be valuable. Given these strategies, an application of Bayes' rule yields the following consistent beliefs:

$$\begin{aligned} \beta^P(p_H, 0) &= \frac{\lambda(1 - q_H)}{\lambda(1 - q_H) + (1 - \lambda)(1 - q_L)} < \lambda, \\ \beta^P(p_H, 1) &= \frac{\lambda q_H}{\lambda q_H + (1 - \lambda)q_L} > \lambda. \end{aligned}$$

Assumptions 1 and 2 now immediately imply that the principal's strategy is sequentially rational. As for the agent, consider type $\alpha_1 = 0$ first. The condition for $s_1^A(p_H, 0) = q_L$ to be sequentially rational is

$$\begin{aligned} u^A(p_H, q_L, 0) + q_L u^A(p_H, q^*(p_H, 0), 0) + (1 - q_L)u^A(p_L, q^*(p_L, 0), 0) \\ \geq u^A(p_H, q_H, 0) + q_H u^A(p_H, q^*(p_H, 0), 0) + (1 - q_H)u^A(p_L, q^*(p_L, 0), 0), \end{aligned}$$

which is satisfied under Assumptions 1 and 4(i). Now consider type $\alpha_1 = a$. The condition for $s_1^A(p_H, a) = q_H$ to be sequentially rational is

$$\begin{aligned} u^A(p_H, q_H, a) + q_H u^A(p_H, q^*(p_H, a), a) + (1 - q_H)u^A(p_L, q^*(p_L, a), a) \\ \geq u^A(p_H, q_L, a) + q_L u^A(p_H, q^*(p_H, a), a) + (1 - q_L)u^A(p_L, q^*(p_L, a), a), \end{aligned}$$

which is satisfied under Assumptions 1 and 4(ii). The resulting expected payoff of the principal in the root of this subgame is

$$\begin{aligned} U_H^P = & \lambda q_H [u^P(p_H, 1) + u^P(p_H, q_H)] \\ & + \lambda(1 - q_H) [u^P(p_H, 0) + u^P(p_L, q_L)] \\ & + (1 - \lambda) q_L [u^P(p_H, 1) + u^P(p_H, q_L)] \\ & + (1 - \lambda)(1 - q_L) [u^P(p_H, 0) + u^P(p_L, q_L)]. \end{aligned}$$

Now consider the principal's choice of the first-period price. The condition $U_H^P \geq U_L^P$ is satisfied under Assumption 4(iii), which implies that $s^P(\emptyset) = p_H$ is sequentially rational for the principal. ■

It remains to be shown that Assumptions 1 – 4 can be satisfied at the same time, so that the low-trust and the high-trust equilibrium coexist. In fact, it can be shown that they are jointly satisfied by a large range of values of the underlying parameters, including values that resemble the payoff structure in the experiment. For instance, let $v = 100$, $c = 15$, $p_L = 20$, $p_H = 40$, $q_L = 1/3$ and $q_H = 2/3$. Also choose $a = 0.4$ and $\lambda = 0.55$. It is easy to show that all assumptions are satisfied by these parameters. It holds that all players' expected payoffs are non-negative in equilibrium (even in each period separately), so that participation constraints would also be satisfied.

Equilibrium Selection

We interpret the historical examples about agent trustworthiness used in the experiment as an equilibrium selection device, so the high-trust equilibrium corresponds to treatment HT and the low-trust equilibrium corresponds to treatment LT.

Consider the low-trust equilibrium first. The offered price and the returned quality are always low on the equilibrium path, in both periods. If the principal actually trembled and mistakenly offered the high price in the first period, both agent types would still respond with low quality. The off-equilibrium behavior thus confirms the pattern shown in the low-trust example.

Now consider the high-trust equilibrium. On the equilibrium path, the first-period price is always high, p_H , while the average second-period price (across many independent repetitions of the game) is given by

$$[\lambda q_H + (1 - \lambda) q_L] p_H + [\lambda(1 - q_H) + (1 - \lambda)(1 - q_L)] p_L,$$

reflecting that the principal offers the high price only if the first-period good turns out to be valuable. Hence, we predict a small endgame effect in prices. As for quality, the average first-period quality is $\lambda q_H + (1 - \lambda)q_L$ in equilibrium. The average second-period quality is $(\lambda q_H)q_H + (1 - \lambda q_H)q_L$, reflecting that only trustworthy agents who produced a good of high value are induced to provide the high quality again in the second period. Hence, there should also be some endgame effect in quality. Notice that a response of quality to price can be observed on the equilibrium path in the high-trust equilibrium. The high price is associated with an average quality of $\lambda q_H + (1 - \lambda)q_L$ in the first period and with an even larger average quality of $\bar{\beta}q_H + (1 - \bar{\beta})q_L$ in the second period. By contrast, the low price in the second period is always associated with the low quality. Similarly, if the principal trembled and offered the low price in the first period, both agent types would respond with low quality. This confirms the responsive pattern where agents choose higher average quality when prices are higher, as shown in the high-trust example.

C. Experimental Instructions

In this section, we provide an English translation of the original German of our main treatment pair HT versus LT and of treatment pair HT-R versus LT-R. Comments in square brackets indicate where the instructions differ between the treatments. The instructions for the treatment pair HT-R2 versus LT-R2 are identical to the respective instructions for HT-R and LT-R, except that they specify that (i) buyers are not obliged to pay the offered price and (ii) actual prices and quality levels are chosen simultaneously. The instructions for treatment HT-Long are identical to the instruction for treatment HT, except that they specify that the game lasts for 25 periods and a slightly different exchange rate between points to account for the fact that subjects were paid for 10 additional periods. The original German instructions for all treatments are available from the authors upon request.

C.1 Buyer Instructions

Instructions for buyers

You are now participating in an economic study. Please read the following instructions carefully. Here you will learn everything you need to know to participate in the study. Please raise your hand if you do not understand something. We will answer your question at your desk.

You will receive an initial endowment of **20 Swiss francs** at the beginning of the study. You can earn additional income during the study by earning **points**. The number of points you earn during the study depends on your choices and on those of the other participants.

All the points you earn during the course of the study will be converted to Swiss francs at the end of the study. The following conversion rate applies:

10 points = 1 Swiss franc

You will receive the monetary amount you earned during the study plus the 20 Swiss francs initial endowment in cash at the end of the study.

The study is divided into individual periods. You must make decisions each period which you enter into the computer. There are a total of 15 periods.

Please note that communication is strictly forbidden during the study. Furthermore, we inform you that you may only use those functions on the computer that are necessary for completing the study. Communication or playing with the computer lead to exclusion from the study. We remain at your disposal to answer any questions you might have.

The 34 participants were divided into two completely independent groups of 17 participants each before the beginning of the study. **You will only interact within your group of 17 participants during the study.** The participants in each group of 17 are then divided into 10 sellers and 7 buyers each.

You are a buyer during the entire study. *[Treatments HT/LT:]* All participants have an identification number that they retain for the entire duration of the study. Your identification number is on the documentation sheet in front of you.

[Treatments HT-R/LT-R:] All participants have an identification number **that changes randomly in each period.** *[Identical instructions from here onwards.]*

Short summary of the procedure of the study

Each buyer can trade a product with a seller in each period of the study. The seller realizes a profit if he/she obtains a sales price that exceeds his/her production costs. The production costs depend on the quality of the product. The higher the quality of the product that the seller chooses, the higher are his/her production costs. The buyer realizes a profit if he/she pays less for the product than its value to him/her. The value of a product is either 100 points (the product is exceptionally good) or 10 points (the product just fulfills its objective). The higher the quality that the seller selects, the higher the probability that the value of the product for the buyer is 100 points.

The study lasts a total of 15 periods. The procedure in an individual period is organized as follows:

1. Each period begins with a **negotiation phase** that lasts for three minutes. Buyers can make purchase offers that sellers can accept during this time.

When making a purchase offer, a buyer must determine **three things**:

- the price he/she offers;
- the quality he/she desires;
- and finally to which seller he/she is directing the purchase offer. Buyers can make two types of purchase offers: private and public. **Private purchase offers** are **only directed to one seller** and can only be accepted by this seller. **Public purchase offers** are directed to **all sellers** and can thus be accepted by any seller.

As a buyer, you can make as many purchase offers as you want in each period. An offer which is made can be accepted at any time. **Each seller and each buyer can only conclude a maximum of one trade in each period.** As there are a total of ten sellers and seven buyers, some sellers will not be able to participate in each period.

2. After the negotiation phase, all sellers who have concluded a trade must determine the product quality that they want to deliver to their buyer. **The seller must not respect the buyer's desired quality.** After all of the sellers have selected their product quality, it will be determined whether the value for each buyer is 100 points or 10 points. The higher the seller's selected quality, the higher the probability that the value for the buyer is 100 points. Once these values have been determined, the earnings for this period are settled for this period. Then the next period begins.

The income from all 15 periods will be added together at the end of the study, converted to Swiss francs, and paid out together with the initial endowment in cash.

Detailed procedure of the study

There are 7 buyers and 10 sellers in your group of 17. You are a **buyer** for the entire study. You will enter your decisions on the computer during the study. The information below shows you in detail how you make your decisions in each period.

1. The negotiation phase

Each period of the study begins with a negotiation phase. Each buyer can conclude a trade with one seller in the negotiation phase. **Each buyer can make as many purchase offers as he/she wants to** during this phase. You will see the following decision screen during each negotiation phase:

Periode			1 von 15			Verbleibende Zeit [sec]: 170		
öffentliche Angebote			Ihre privaten Angebote			Ihre Identifikationsnummer		
Käufer	Preis	gew. Q.	Preis	gew. Q.	an Verk.	Hier machen Sie Ihre Angebote		
						☐ öffentlich ☐ privat falls privat, an welchen Verkäufer?		
						Ihr Preis		
						gewünschte Qualität		
						OK		
						☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 ☐ 8 ☐ 9 ☐ 10		
						Ihr Verkäufer Ihr Preis Ihre gew. Q.		

[Text on screen: Period 1 of 15 / remaining time / public purchase offers buyer price desired quality / your private offers price desired quality to seller / your identification number / you make your offers here / public / private / if private, to which seller? / your price / desired quality / OK / your seller / your price / your desired quality]

- You see which period you are in at the upper left corner of the screen. The remaining time in this negotiation phase appears at the upper right corner. **The negotiation phase lasts three minutes** (=180 seconds) in each negotiation phase. Once the time has expired, the negotiation phase is over. Further purchase offers may neither be made nor accepted in this period.

[Only treatments HT-R/LT-R, but not HT/LT, include the following bullet point:]

- The next item you see is your identification number. **This identification number is randomly redetermined in each period.** This applies to all study participants, i.e. for all sellers and for all buyers.
- As soon as you see the screen above, the negotiation phase is opened. You as a buyer now have the opportunity to make purchase offers to the sellers. To do this you must determine three things on the right side of the screen:

a) You first must determine if you want to make a public or private purchase offer:

- **Public purchase offers**

Public purchase offers are notified to all participants in the market. All sellers see all public purchase offers on their screens. **Each seller** can thus accept a public purchase offer. You as buyer also see all public purchase offers from the other buyers.

If you want to make a public purchase offer, click on the field "public" with your mouse.

- **Private purchase offers**

Private purchase offers are only directed towards **one seller**. Only this seller learns of the offer, and only this seller may accept the purchase offer. No other sellers and buyers in the market will learn about this offer.

If you want to make a private purchase offer, click on the field "private" with the mouse. Then indicate in the field below to which seller you direct the purchase offer. All ten sellers have an identification number (seller 1, seller 2, ..., seller 10). [Treatments HT/LT:] The sellers retain this number for the duration of the study. [Treatments HT-R/LT-R:] **This identification number varies randomly in each period of the study.** [Identical instructions from here onwards.] To direct an offer to a specific seller, enter the seller's number (e.g. "4" for seller 4).

b) After you have determined to whom you want to direct your purchase offer, you must determine your **purchase offer**. Enter this in the field "your price". The purchase offer may neither be less than 0 nor greater than 100.

$$0 \leq \text{purchase offer} \leq 100$$

c) Finally, you must then enter the product quality you desire. Enter this in the field "desired quality". The **desired product quality** may neither be less than 1 nor higher than 9:

$$1 \leq \text{desired product quality} \leq 9$$

- After you have completely determined your purchase offer, you must click on the "OK" button to publicize the offer. You may revise your offer until you click the "OK" button. After you click on the "OK" button, your purchase offer will appear to all sellers to whom it was directed.

- You see the heading "public offers" on the left side of your screen. All public offers in the current negotiation phase appear here. Both your own offers as well as the public offers from the other buyers appear here. You can see which buyer made the offer, the price he/she offers, and the quality he/she desires. *[Treatments HT/LT:]* All buyers also have an identification number in the study that applies for the entire duration of the study (buyer 1, buyer 2, ... buyer 7). *[Treatments HT-R/LT-R:]* All buyers in the study also have an identification number between one and seven in each period of the study (buyer 1, buyer 2, ... buyer 7). **The buyers' identification number is also randomly redetermined in every period of the study.** *[Identical instructions from here onwards.]*
- The private offers you made in the current negotiation phase are listed under the heading "your private offers" in the middle of the screen. Here you see to which sellers you made offers, which prices you offered in each case, and which quality you desired.
- **Each buyer can make as many private and public offers as he/she wants in each period.** Every purchase offer you make can be accepted at any time during the negotiation phase.
- **Each buyer can only conclude one trade in each period.** As soon as one of your purchase offers is accepted, you will be informed which seller accepted your offer. The number of the seller who accepted the offer, your price offer and your desired quality appear in the lower right corner of your screen. As you can only conclude one trade per period, your other purchase offers will be automatically deleted at this time. Furthermore, you cannot make any other purchase offers in this period.
- **Each seller can only conclude one trade at most in each period.** You will be continuously informed about which sellers have not yet concluded a trade. You see ten fields at the lower corner of your screen. Once a seller has accepted a purchase agreement, an "X" appears in the box before his/her identification number. You can no longer make a private offer to those sellers who have already accepted an offer.
- The negotiation period is over as soon as all seven buyers have concluded a trade or the three minutes are up.
- No buyer is forced to make a purchase offer, and no seller is forced to accept an offer.

2. Determination of the actual product quality

Once the negotiation phase is over, all sellers who have concluded a trade must decide which product quality they want to deliver to their buyer. **The product quality that you desired in your offer is not binding for your seller.** Your seller can select exactly the product quality you desired, but can also choose a higher or lower product quality. The quality your seller selects must be an integer between 1 and 9.

$$1 \leq \text{actual product quality} \leq 9$$

While your seller selects the actual product quality, we ask that you indicate on a separate screen the actual product quality that you expect. We also ask you to state how confident you are about your estimate.

How are the incomes calculated?

Your income:

- If you do not conclude a trade during the period in question, you will earn an income of 0 points in the period.
- If one of your offers was accepted, your income depends on the price you offer and whether a product value of 10 or 100 points is determined. **The higher the quality the seller selects, the higher is the probability that the value of your product will be 100.** Your income is calculated as follows:

$$\begin{aligned} \text{Your income} &= 100 - \text{price, if the high product value is determined} \\ \text{Your income} &= 10 - \text{price if the low product value is determined} \end{aligned}$$

Your expected income is thus higher, if the product quality your seller delivers is high. At the same time, your income is higher, if the price that you must pay for the product is lower.

The probability for a value of 100 points depends on the selected quality as follows:

Quality	1	2	3	4	5	6	7	8	9
Probability of a value of 100	10%	20%	30%	40%	50%	60%	70%	80%	90%

If the seller selects a quality of 1, the probability that the product will have a value of 100 for the buyer is 10%. The probability is 20% for a quality of 2, and so on. The probability is 90% for the maximum quality of 9.

Your seller's income:

- If a seller does not conclude a trade in the negotiation phase, he/she earns the income of 5 points in the corresponding period.
- If a seller accepts a purchase offer, his/her income equals his/her earned price less the production costs he/she incurs. Your seller's income is calculated as follows:

$$\text{Your seller's income} = \text{price less production costs}$$

The higher the seller's selected quality, the higher are his/her production costs. The production costs for each product quality are listed in the following table:

Quality	1	2	3	4	5	6	7	8	9
Production costs	0	1	2	4	6	8	10	12	15

The higher the price, the higher is your seller's income. Furthermore, the lower the product quality he delivers, the higher his/her income.

The sellers' and buyers' incomes are all calculated in the same manner. **Each seller learns whether the realized value of a product is 10 or 100 points, and can thus calculate his/her seller's income. However, a buyer can only guess his seller's income, since he/she cannot observe the selected quality. A buyer can only see which product value is realized.** [The following sentence is included in treatments HT/LT only.] In each period, both buyers and sellers will learn of their trade partner's identification number (ID).

Please note that buyers and sellers can also incur a loss in any period. This must be paid from your initial endowment or from income earned in other periods.

You will learn of your income in a **profit screen** (see next page). The following information will be notified to you there:

- The seller with whom you concluded a trade.
- The price you offered.
- The quality you requested.
- Whether a product value of 10 or 100 was determined.
- The income you earned in this period.
- You only know that your seller's income results from the price less production costs. Since you cannot see which quality was actually chosen, you do not know your seller's production costs.

Periode 1 von 15	Verbleibende Zeit [sec]: 30
Ihre Identifikationsnummer Sie haben folgenden Handel abgeschlossen ID Ihres Verkäufers Preis gewünschte Qualität realisierter Wert des Produkts Ihr Einkommen = weiter	

[Text on screen: Period 1 of 15 / Remaining time [sec] / Your identification number / You concluded the following trade / Your seller's ID / Price / Requested quality / Realized value of the product / Your income = / Continue]

Please enter all the information in the enclosed documentation sheet. Once the profit screen disappears, the period is over. The negotiation phase of the next period then begins. *[The following sentence is included in treatments HT-R/LT-R only.]* You will receive a new, randomly determined identification number for the next period, as will all other participants in the study. Once you have finished looking at the profit screen, please press the "continue" button.

The sellers also have a profit screen where they are also informed about the information above. The sellers see their buyers' ID, the price, the requested and actual product quality, whether the realized product value is 10 or 100 points, and the income for the seller and buyer that results.

The study does not begin until all participants are completely informed about the study procedure. In order to confirm this, we ask that you solve a few practice questions.


Furthermore, we will conduct two **test periods of the negotiation phase** so that you can become more familiar with the computer. These test periods will not be included in the final result and will not be paid out. After the test periods, the study will begin, lasting for a total of 15 periods.

Example


Before the present study begins, we would like to inform you of the study results that we observed in past sessions. You can use this information when you make your decisions today.

The **bars** in the left chart show the average **quality** that the sellers chose for the various prices.

The left diagram shows that a seller on average selected an actual quality of 1.1 for a price offer up to 9. By offers from 10 up to and including 19, quality of 2.6 was selected, for the range from 20 to 29 a quality of 3.6, and for prices from 30 to 39 a quality of 5.3. An average quality of 8.2 was selected for prices of 40 or higher. You can clearly see that the selected quality increases on average if the price increases.



Sellers' quality selection



Buyers' income

The **line** in the graph on the right emphasizes the correlation between price and **income** that is realized with these quality values.

For prices up to 9, the quality is only 1.1, but a lower price is also paid. On average, a buyer's income is approximately 13 points. If prices from 20 to 29 are paid and the average quality is 3.6, an income of 19 results. Although the price paid is considerably higher, a total higher income results since the average quality is also considerably higher. For prices of 40 or more, a buyer realizes an average income of 34 points.


[Treatments LT/LT-R:]

Example


Before the present study begins, we would like to inform you of the study results that we observed in past sessions. You can use this information when you make your decisions today.

The **bars** in the left chart show the average **quality** that the sellers chose for the various prices.

The left diagram shows that a seller on average selected an actual quality of 1.4 for a price offer up to 9. By offers from 10 up to and including 19, quality of 1.6 was selected, for the range from 20 to 29 a quality of 2.8, and for prices from 30 to 39 a quality of 3.1. An average quality of 1 was selected for prices of 40 or higher. You can see that the selected quality barely increases on average if the price increases.



Sellers' quality selection



Buyers' income

The **line** in the graph on the right emphasizes the correlation between price and **income** that is realized with these quality values.

For prices up to 9, the quality is only 1.4, but a lower price is also paid. On average, a buyer's income is approximately 14 points. If, for example, prices from 20 to 29 are paid and the average quality is 2.8, an income of 12 results. Although the price paid is somewhat higher, a total lower income results since the price paid is considerably higher. For prices of 40 or more, a buyer incurs a loss in excess of 20 points.

Practice questions

Please solve these questions completely and **show how you reached your answer**. If you have questions, please raise your hand. False answers have no consequence for your payment at the end of the study.

Question 1

You did not make a purchase offer in a period. How high is your income in the period?

Your income =

Question 2:

Your purchase offer with a price of 30 and a requested quality of 9 is accepted.

Your income when the value of the product is 100 equals =

Your income when the value of the product is 10 equals =

Your seller's income =

How great is the probability that the value of the product for you is 100 points if the seller selects an actual quality of 8?

How high is your seller's income in this case?

Question 3:

Your purchase offer with a price of 60 and a requested quality of 9 is accepted.

Your income when the value of the product is 100 equals =

Your income when the value of the product is 10 equals =

How great is the probability that the value of the product for you is 100 points if the seller selects an actual quality of 6?

How high is your seller's income in this case?

Question 4:

Your purchase offer with a price of 10 and a requested quality of 2 is accepted.

Your income when the value of the product is 100 equals =

Your income when the value of the product is 10 equals =

How great is the probability that the value of the product for you is 100 points if the seller selects an actual quality of 5?

How high is your seller's income in this case?

Question 5:

Your purchase offer with a price of 10 and a requested quality of 6 is accepted.

Your income when the value of the product is 100 equals =

Your income when the value of the product is 10 equals =

How great is the probability that the value of the product for you is 100 points if the seller selects an actual quality of 2?

How high is your seller's income in this case?

Question 6

A seller did not accept a purchase offer in a negotiation phase. How high is this seller's income in this period?

Seller's income =

Question 7

You made several purchase offers in a negotiation phase. None of these offers were accepted by a seller. How high is your income in this period?

Your income =

Question 8:

Look at the example on page 9 of these instructions that provides information about a past session of this study.

How did the seller's income change when the offered price increased?

☐ The income increased.

☐ The income decreased

Explain briefly the reason for this correlation between a buyer's income and the price offered.

Please raise your hand when you have solved these practice questions. We will then come to your seat and check your answers.

C.2 Seller Instructions

Instructions for sellers

You are now participating in an economic study. Please read the following instructions carefully. Here you will learn everything you need to know to participate in the study. Please raise your hand if you do not understand something. We will answer your question at your desk.

You will receive an initial endowment of **20 Swiss francs** at the beginning of the study. You can earn additional income during the study by earning **points**. The number of points you earn during the study depends on your choices and on those of the other participants.

All the points you earn during the course of the study will be converted to Swiss francs at the end of the study. The following conversion rate applies:

10 points = 1 Swiss franc

You will receive the monetary amount you earned during the study plus the 20 Swiss francs initial endowment in cash at the end of the study.

The study is divided into individual periods. You must make decisions each period which you enter into the computer. There are a total of 15 periods.

Please note that communication is strictly forbidden during the study. Furthermore, we inform you that you may only use those functions on the computer that are necessary for completing the study. Communication or playing with the computer lead to exclusion from the study. We remain at your disposal to answer any questions you might have.

The 34 participants were divided into two completely independent groups of 17 participants each before the beginning of the study. **You will only interact within your group of 17 participants during the study.** The participants in each group of 17 are then divided into 10 sellers and 7 buyers each.

You are a seller during the entire study. *[Treatments HT/LT:]* All participants have an identification number that they retain for the entire duration of the study. Your identification number is on the documentation sheet in front of you.

[Treatments HT-R/LT-R:] All participants have an identification number **that changes randomly in each period.** *[Identical instructions from here onwards.]*

Short summary of the procedure of the study

Each buyer can trade a product with a seller in each period of the study. The seller realizes a profit if he/she obtains a sales price that exceeds his/her production costs. The production costs depend on the quality of the product. The higher the quality of the product that the seller chooses, the higher are his/her production costs. The buyer realizes a profit if he/she pays less for the product than its value to him/her. The value of a product is either 100 points (the product is exceptionally good) or 10 points (the product just fulfills its objective). The higher the quality that the seller selects, the higher the probability that the value of the product for the buyer is 100 points.

The study lasts a total of 15 periods. The procedure in an individual period is organized as follows:

1. Each period begins with a **negotiation phase** that lasts for three minutes. Buyers can make purchase offers that sellers can accept during this time.

When making a purchase offer, a buyer must determine **three things**:

- the price he/she offers;
- the quality he/she desires;
- and finally to which seller he/she is directing the purchase offer. Buyers can make two types of purchase offers: private and public. **Private purchase offers** are **only directed to one seller** and can only be accepted by this seller. **Public purchase offers** are directed to **all sellers** and can thus be accepted by any seller.

A buyer can make as many purchase offers as he/she wants in each period. **You as a seller can only accept one purchase offer at most per period.** If you accept a seller's offer, you conclude a trade with this seller for this period. **Buyers, too, can only conclude one trade at most in a period.** As there are a total of ten sellers and seven buyers, some sellers will not be able to participate in each period.

2. After the negotiation phase, all sellers who have concluded a trade must determine the product quality that they want to deliver to their buyer. **As a seller, you must not respect the buyer's desired quality.** After all of the sellers have selected their product quality, it will be determined whether the value for each buyer is 100 points or 10 points. The higher the seller's selected quality, the higher the probability that the value for the buyer is 100 points. Once these values have been determined, the earnings for this period are settled for this period. Then the next period begins.

The income from all 15 periods will be added together at the end of the study, converted to Swiss francs, and paid out together with the initial endowment in cash.

Detailed procedure of the study

There are 7 buyers and 10 sellers in your group of 17. You are a **seller** for the entire study. You will enter your decisions on the computer during the study. The information below shows you in detail how you make your decisions in each period.

1. The negotiation phase

Each period of the study begins with a negotiation phase. Each buyer can conclude a trade with one seller in the negotiation phase. To do this, buyers can make purchase offers to sellers. As a seller, you can accept one of the offers made to you in each period. You will see the following decision screen during each negotiation phase:

Periode 1 von 51		Verbleibende Zeit [sec]: 116
Ihre Identifikationsnummer		
Private Angebote an Sie		öffentliche Angebote
von Käufer	Preis	gew. Q.
akzeptieren		akzeptieren
Ihr Käufer		Ihr Preis
gew. Q.		

[Text on screen: Period 1 of 15 / remaining time / your identification number / private offers to you / from seller / price / desired quality / public offers / from seller / price / desired quality / accept / your seller / your price / desired quality]

- You see which period you are in at the upper left corner of the screen. The remaining time in this negotiation phase appears in seconds at the upper right corner. **The negotiation phase lasts three minutes** (=180 seconds) in each negotiation phase. Once the time has expired, the negotiation phase is over. Further purchase offers may neither be made nor accepted in this period.

[Only treatments HT-R/LT-R, but not HT/LT, include the following bullet point:]

- The next item you see is your identification number. **This identification number is randomly redetermined in each period.** This applies to all study participants, i.e. for all sellers and for all buyers.
- As soon as you see the screen above, the negotiation phase is opened. You as a seller now have the opportunity to accept purchase offers the sellers make to you. There are two types of purchase offers that you can accept:

- **Private purchase offers to you**

Each buyer may offer you private purchase offers. These offers **are only made to you**, and **only you can accept them**. No other sellers or buyers will learn about this offer. If you receive private offers, they appear on the left side of your screen under the title "private offers to you". The buyer's offer contains the following information: the identification number of the buyer who is making the offer, the price he/she offers for the product, and the quality he/she desires. If you want to accept a private offer, first click on the line where the private offer is entered with the mouse. The corresponding offer will then be marked. If you want to definitely accept the offer, click on the "accept" button at the lower right corner. You can change your selection up until you click on the "accept" button.

- **Public purchase offers**

Each buyer can make public purchase offers. Public purchase offers are notified to all sellers. All sellers see all public purchase offers on their screens. **Each seller can thus accept a public purchase offer.** If a seller makes a public offer, it appears on the right side of your screen under the title "public offers". The offer again consists of the identification number of the buyer who is making the offer, the price he/she offers for the product, and the quality he/she desires. All other sellers and buyers receive this information. If you want to accept a public offer, the same procedure as for the private offers applies. First click on the line where the offer is entered. If you want to accept the offer definitely, click on the "accept" button on the lower right side. You can change your selection up until you click on the "accept" button.

- As soon as you click on the "accept" button, you will see which offer you accepted on the bottom line of your screen.
- **Each seller can only conclude one trade in a period.** Once you have accepted a purchase offer, you cannot accept any more offers.

All buyers must accept the following rules for their purchase offers:

- The buyer's offer may not be less than 0 and may not be greater than 100:

$$0 \leq \text{purchase offer} \leq 100$$

- A buyer's **desired product quality** cannot be less than 1 and cannot be greater than 9.

$$1 \leq \text{desired product quality} \leq 9$$

- **Each buyer can make as many private and public purchase offers in each period as he/she wants to.** Each purchase offer a buyer makes in a period can be accepted during the negotiation phase.
- **Each buyer can only conclude one trade in each period.** As soon as a buyer's purchase order is accepted, he/she will learn which seller accepted this offer. Since each buyer can only conclude one trade in each period, his/her remaining purchase offers are automatically deleted. He/She also cannot make any other purchase offers.
- The negotiation period is over as soon as all seven buyers have concluded a trade or the three minutes are up.
- No buyer is forced to make a purchase offer, and no seller is forced to accept an offer.

2. Determination of the actual product quality

Once the negotiation phase is over, all sellers who have concluded a trade must decide which product quality they want to deliver to their buyers. **The product quality that your buyer desired in your offer is not binding for you as seller.** You can select exactly the product quality your buyer desired, but you can also choose a higher or lower product quality. If you concluded a trade in a negotiation phase, the entry of the product quality appears on the following screen:

Periode 1 von 15 Verbleibende Zeit [sec]: 24

Sie haben folgendes Angebot akzeptiert

von Käufer

Preis

gewünschte Qualität

Bestimmen Sie die tatsächliche Qualität

OK

[Text on screen: Period 1 of 15 / remaining time [sec] / You accepted the following offer / from buyer / price / desired quality / determine the actual quality / OK]

- In order to select your actual product quality, enter the value for the quality in the field "determine the actual quality" and click on the "OK" button. You can change your selection until you click on the "OK" button.

The quality you select must be an integer between one and nine.

$$1 \leq \text{actual product quality} \leq 9$$

How are the incomes calculated?

Your income:

- If you **do not conclude a trade** in the negotiation phase, you will earn an income of **5 points** in the corresponding period.
- If you accept a purchase offer, your income depends on the accepted purchase and the product quality you select. Your income is calculated as follows:

$$\text{Your income} = \text{price less production costs}$$

- The higher your quality, the higher are your production costs. The production costs for each product quality are listed in the following table:

Quality	1	2	3	4	5	6	7	8	9
Production costs	0	1	2	4	6	8	10	12	15

- The lower your selected quality, the higher is your income. Furthermore, the higher your buyer's offered price, the higher is your income.

Your buyer's income:

- If a buyer does not conclude a trade during the period in question, he/she will earn an income of 0 points in the period.
- If one of his/her offers was accepted, his/her income depends on the price he/she offers and whether a product value of 10 or 100 points is determined. **The higher the quality the seller selects, the higher is the probability that the value of the product will be 100.** Your buyer's income is calculated as follows:

$$\begin{aligned}\text{Your income} &= 100 - \text{price, if the high product value is determined} \\ \text{Your income} &= 10 - \text{price if the low product value is determined}\end{aligned}$$

Your buyer's expected income is thus higher, if the product quality you deliver is high, as a higher quality increases the probability that the product value will be 100. At the same time, his/her income is higher, if the price that he/she must pay for the product is lower.

The probability for a value of 100 points depends on the selected quality as follows:

Quality	1	2	3	4	5	6	7	8	9
Probability of a value of 100	10%	20%	30%	40%	50%	60%	70%	80%	90%

If a quality of 1 is selected, the probability that the product will have a value of 100 for the buyer is 10%. The probability is 20% for a quality of 2, and so on. The probability is 90% for the maximum quality of 9.

The sellers' and buyers' incomes are all calculated in the same manner. **Each seller learns whether the realized value of a product is 10 or 100 points, and can thus calculate his/her realized income. However, a buyer can only guess his seller's income, since he/she cannot observe the selected quality. A buyer can only see which product value is realized.** *[The following sentence is included in treatments HT/LT only.]* In each period, both buyers and sellers will learn of their trade partner's identification number (ID).

Please note that buyers and sellers can also incur a loss in any period. This must be paid from your initial endowment or from income earned in other periods.

You will learn of your income in a **profit screen** (see next page). The following information will be notified to you there:

- The buyer with whom you concluded a trade.
- The price the buyer offered.
- The quality the buyer desired.
- The actual quality that you selected.
- Whether a product value of 10 or 100 was determined.
- The income your buyer earned in this period.
- The income you earned in this period.

Periode 1 von 15	Verbleibende Zeit [sec]: 34
--	-----------------------------

Ihre Identifikationsnummer

Sie haben folgenden Handel abgeschlossen

ID Ihres Käufers

Preis

gewünschte Qualität

tatsächlich gewählte Qualität

realisierter Wert des Produkts

Einkommen Ihres Käufers =

Ihr Einkommen =

weiter

[Text on screen: Period 1 of 15 / remaining time / your identification number / You concluded the following trade / your buyer's ID / price/ desired quality / actually selected quality / realized value of the product / your buyer's income / your income / continue]

Please enter all the information in the enclosed documentation sheet. Once the profit screen disappears, the period is over. The negotiation phase of the next period then begins. *[The following sentence is included in treatments HT-R/LT-R only.]* You will receive a new, randomly determined identification number for the next period, as will all other participants in the study. Once you have finished looking at the profit screen, please press the "continue" button.

The buyers also have a profit screen where they are also informed about the information past period. The buyers see their seller's ID, the price, and the desired product quality. However, the buyers cannot observe the product quality that was actually chosen. The buyers only see if a product value of 10 or 100 was realized.

The study will not begin until all participants are completely informed about the study procedure. In order to confirm this, we ask that you solve a few practice questions.

Furthermore, we will conduct two **test periods of the negotiation phase** so that you can become more familiar with the computer. These test periods will not be included in the final result and will not be paid out. After the test periods, the study will begin, lasting for a total of 15 periods.

Practice questions

Please solve these questions completely and **show how you reached your answer**. If you have questions, please raise your hand. False answers have no consequence for your payment at the end of the study.

Question 1

You did not accept a purchase offer in a period. How high is your income in the period?

Your income =

Question 2:

You accepted a purchase offer with a price of 60 and a desired quality of 9. You select the actual quality of 9.

Your income =

Your buyer's income when the value of the product is 100 equals =

Your buyer's income when the value of the product is 10 equals =

How great is the probability that the value of the product for you is 100 points if the seller selects an actual quality of 9?

Question 3:

You have accepted purchase offer with a price of 60 and a desired quality of 9. You choose an actual quality of 4.

Your income =

Your buyer's income when the value of the product is 100 equals =

Your buyer's income when the value of the product is 10 equals =

How great is the probability that the value of the product for you is 100 points if the seller selects an actual quality of 4?

Question 4:

You have accepted purchase offer with a price of 40 and a desired quality of 2. You choose an actual quality of 5.

Your income =

Your buyer's income when the value of the product is 100 equals =

Your buyer's income when the value of the product is 10 equals =

How great is the probability that the value of the product for you is 100 points if the seller selects an actual quality of 5?

Question 5:

You have accepted purchase offer with a price of 30 and a desired quality of 6. You choose an actual quality of 6.

Your income =

Your buyer's income when the value of the product is 100 equals =

Your buyer's income when the value of the product is 10 equals =

How great is the probability that the value of the product for you is 100 points if the seller selects an actual quality of 6?

Question 6

A buyer made several purchase offers in a negotiation phase. None of these offers were accepted by a seller. How high is the buyer's income in the period in question?

Buyer's income =

Please raise your hand when you have solved these practice questions. We will then come to your seat and check your answers.