



Working Papers

www.cesifo.org/wp

Cooperation and Mistrust in Relational Contracts

Holger Herz
Armin Schmutzler
André Volk

CESIFO WORKING PAPER NO. 6045

CATEGORY 13: BEHAVIOURAL ECONOMICS

AUGUST 2016

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

ISSN 2364-1428

CESifo

Center for Economic Studies & Ifo Institute

Cooperation and Mistrust in Relational Contracts

Abstract

Work and trade relationships are often governed by relational contracts, in which incentives for cooperative action today stem from the prospective future benefits of the relationship. In this paper, we study how reductions in clarity about the financial consequences of actions, induced by incomplete information about the costs of providing quality, affect relational contracts in buyer-seller relationships. Under incomplete information, payoffs to actions become private information. This can impede the joint understanding of what constitutes cooperative behavior, and may thus inject mistrust into relationships, even if credibility is held constant. Comparing seller-buyer relationships with and without complete information about seller costs in the laboratory, we find that such a lack of clarity has effects on the terms of relational contracts. However, these effects only concern the distribution of rents, and not efficiency.

JEL-Codes: D010, D030, L140, L200.

Keywords: relational contracts, incomplete information, experiments.

Holger Herz
Department of Economics
University of Fribourg
Boulevard des Pérolles 90
Switzerland – 1700 Fribourg
holger.herz@unifr.ch

Armin Schmutzler
Department of Economics
University of Zurich
Blümlisalpstrasse 10
Switzerland – 8006 Zurich
armin.schmutzler@econ.uzh.ch

André Volk
Department of Economics
University of Zurich
Blümlisalpstrasse 10
Switzerland – 8006 Zurich
andre.volk@econ.uzh.ch

August 15, 2016

The paper benefited from helpful comments by Claudia Geiser, Manuel Grieder, Shuo Liu, Michel Marechal, Michael Powell, Florian Zimmermann and seminar audiences at Düsseldorf (DICE), Lausanne, Madrid, Magdeburg, Nottingham Business School, WHU Koblenz, Zürich, the Workshop on Relational Contracts at Ammersee, and the German Economic Association Meeting in Hamburg. We gratefully acknowledge financial support of the Richard-Büchner-Stiftung. The project received the approval of the IRB at the faculty of economics, business administration, and computer science of the University of Zurich.

1 Introduction

The theory of relational contracts investigates the conditions under which cooperative and efficient behavior is sustainable through repeated interaction even if contracts are not enforceable through third parties. If the individual benefits of upholding the relational contract are sufficiently large for each trading party, and this is common knowledge, then it is in everybody's self-interest to act cooperatively. Thus the terms of the relational contract simply need to ensure that each party individually benefits more from upholding the relationship than from reneging on the relational contract. In this sense, sustaining relational contracts is a matter of credibility.

According to this view, maintaining relational contracts should be relatively simple whenever the gains from trade are sufficiently high. However, Gibbons and Henderson (2012; 2013) have recently argued that the sole focus on credibility may be insufficient, because efficient relational contracts need to solve the “twin problem of credibility and clarity” (Gibbons and Henderson (2012), p. 1350). In addition to credibility, a mutual understanding of the terms of the relational contract is required: The parties need to know which actions are in the spirit of the contract and which actions violate it. Such clarity is hard to generate, which might be one reason why seemingly similar enterprises display persistent performance differences (Syverson, 2011). One important aspect of clarity is so-called “relational knowledge”, which encompasses “the payoff to cooperation for each party, of each party's ability and incentive to defect, and of the actions and payoffs that constitute punishment” (Gibbons and Henderson (2012), p. 1355).

This paper is, to our knowledge, the first to study how variations in relational knowledge affect the terms and efficiency of relational contracts. We conducted a series of experiments, with a total of 496 subjects, in which a seller and a buyer are matched in pairs and repeatedly engage in the transaction of a good of heterogeneous quality that is structured via a *trust contract*. Higher quality is valuable to the buyer, but costly to the seller. In the first stage of each transaction, the buyer pays a price upfront. In the second stage of each transaction, the seller chooses the quality q of the good. Quality is chosen from $q \in \{1, \dots, 10\}$, and the value of the good to the buyer is given by $10q$. The price is chosen from $p \in \{0, \dots, 100\}$. Neither quality nor price are contractible. There are two seller types, which differ in their cost of providing quality. For the high cost type, the absolute and marginal costs of providing quality are always larger, except for the lowest quality level. Each pair engaged in 15 transactions. Relational knowledge about the payoff from transacting is manipulated by exogenously varying the information about the seller's true costs: In the complete information treatments, seller costs are common knowledge, whereas in the incomplete information treatments, they are private information. However, at the beginning of the latter treatments, the seller can send a cheap talk

cost signal to the buyer. Consequently, the buyer only knows the ex-ante probabilities of high or low costs and observes the non-verifiable cost signal.

Relational knowledge is therefore reduced in the incomplete information treatment, since the buyer cannot know the seller's payoff with certainty. However, credibility cannot be lower than in the case of complete information with high costs. For any given combination of price and quality under incomplete information, the seller has lower profits if his costs are high than if they are low. It is therefore intuitively clear that, no matter what the true type of the seller and his cost signal are, the incentives to adhere to the terms of the relationship are never lower than for the high cost type under complete information. If the seller's true costs are indeed high, credibility is unaltered relative to the high cost complete information case. If the seller's true costs are low, the seller has even more to gain from adhering to the terms of the contracts since seller profits are higher. In this sense, in comparison to high cost complete information relations, credibility is never adversely affected by incomplete information.¹

Why then should the reduction in relational knowledge impede building relational contracts? It has been argued that mutual trust is crucial for upholding relationships (Zaheer and Venkatraman, 1995; Bachmann and Zaheer, 2006). Moreover, evidence suggests that dishonest actions and attempts to gain (what are perceived to be) unfair advantages through acts of cheating frequently meet with resistance and punishment (Brandts and Charness, 2003). In the incomplete information treatment, a trading partner may therefore interpret high cost signals as potentially dishonest attempts by actual low cost sellers to gain such advantages, and he may therefore doubt their truthfulness.

Such doubts can be detrimental to building relational contracts because under incomplete information, it is not clear what the fair terms of the relationship are. When deciding on how to distribute the gains from trade, should the cost signal be taken at face value? Or can trading partners agree that rent distribution is based on expected rather than signaled costs? If so, what is the appropriate expectation? In such an environment, if one trading partner takes an action that he believes to be in the spirit of the relational contract, the other one may nevertheless regard it as violating the relational contract. Doubts about the truthfulness of the seller and ambiguity regarding the fair terms of the relational contract may therefore translate into behavior. This may endanger the establishment of efficient relational contracts even though credibility is unaltered. To see this logic more clearly, consider the following intuitive reasoning:

- (S1) When buyers believe that a seller has high costs, they are prepared to compensate quality increases with higher price increases.

¹In order to get a clean identification of the effect of reduced clarity, keeping credibility unaltered, our analysis will therefore focus on actual high cost relationships only. In Section 3, we provide theoretical foundations for the intuitive argument presented here.

- (S2) However, when buyers see a high signal under incomplete information, they doubt that it is truthful.
- (S3) Thus, a buyer might compensate a seller who signals high cost less for higher quality than she would under complete information.
- (S4) This can trigger a quality reduction by high-cost sellers: Even if they understand the informational constraints of the buyer, they may think they do not receive sufficient compensation for their quality.

If steps (S1)-(S4) hold, high cost relationships with incomplete information may be harder to sustain because mutual understanding about what constitutes a cooperative action in the relationship is harder to achieve.

Turning to our data, we find that the average quality provided in high cost relationships is 5.4 under complete information, whereas it is 5.6 in the incomplete information treatments. Therefore, contradicting the hypothesis that reductions in relational knowledge reduce efficiency, the quality is even slightly higher with incomplete information than without, but insignificantly so. Moreover, the development of average quality over time is almost identical between treatments. Consequently, we find no statistically significant or economically meaningful difference in quality: High cost relationships turn out to be equally efficient under complete and incomplete information. To fully understand the mechanisms and intuition underlying this finding, we need to examine the data more closely:

First, focusing on the complete information treatment only, we find that high cost sellers receive a higher price for the same quality than low cost sellers, which is consistent with (S1). Buyers are indeed willing to compensate sellers for higher costs. Consequently, the terms of trade that are deemed acceptable and the resulting distribution of surplus crucially depend on relational knowledge regarding the sellers true costs and true profits from trade.

Second, consistent with (S2), the cost signal is not truthful in general, though it has some informational value: Almost without exception, the high cost sellers state their type correctly. Two thirds of the low cost sellers, however, also claim to have high costs. Moreover, the buyers' beliefs reveal that they anticipate this behavior, so that doubts about true seller types indeed exist following high cost signals. Thus, under incomplete information, buyers do not know the true cost type of sellers following high cost signals and hence the proper compensation of a seller who signals high costs is ambiguous. In particular, true high cost sellers cannot credibly signal that they indeed have high costs.

Indeed, when buyers observe a high cost signal in the incomplete information treatment, they indicate significantly higher desired quality for a given price in the first

period, compared to the high cost complete information case. It hence appears that they desire altered contractual terms. Nonetheless, step (S3) breaks down: Despite the buyers' doubts and higher demanded quality, the empirical relationship between price and quality in actual high cost relationships under complete and incomplete information is not significantly different. This implies that buyers end up compensating sellers as if the cost signal was truthful, i. e., doubts about seller costs do not lead to lower prices for a given quality. In turn, the expected downward spiral for efficiency (S4) does not materialize.

Taken together, it appears that buyers and sellers quickly coordinate on taking the cost signal at face value, despite the buyer's doubts about its truthfulness. This enables the pairs to build and maintain relational contracts on the same terms as under high cost complete information, consistent with the view that maintaining relational contracts is purely a matter of credibility.

Another interesting, and potentially crucial finding in our data is that sellers capture a significantly larger share of the net surplus, realizing 20-50% higher profits than buyers. One possible cause for this asymmetry might be the order of moves within a transaction. Since the seller moves second, his quality choice effectively controls the distribution of rents from the transaction, which may enable him to self-servingly shift the terms of the contract towards a price-quality combination that is favourable to the seller.

If moving second in a transaction is indeed decisive for the terms of the contract, then the structure of moves within the transaction might be a crucial reason why buyers' doubts about the honesty of sellers did not affect actions. The fact that sellers closed the transaction through their choice of quality enabled them to always act in ways consistent with their cost signal. Buyers, on the other hand, lacked the capability to effectively control the distribution of rents. Given that they had to pay the price upfront, buyers may have had no choice but to give the sellers the benefit of the doubt and to accept contractual terms that take the cost signal at face value.

To understand whether this mechanism is indeed present, we conducted a second experiment in which we switched the order of moves within the transaction, and structured the transaction via a *bonus contract*. The seller moves first, chooses quality and indicates a desired price. The buyer then moves second and chooses a price. All other aspects of the experiment remain unchanged. The altered order of moves enables the buyer to determine the distribution of the gains from trade through the price choice, enabling him to translate concerns about honesty into action.

Our second experiment indeed reveals significant effects of incomplete information on the terms of the ensuing relational contracts in high cost relationships. Under incomplete information, prices react less to quality, i. e., buyers are not willing to compensate

sellers in the same way as in the complete information benchmark. When buyers move second, they translate their doubts about the truthfulness of the cost signal into action and are no longer willing to pay prices as if the cost signal was truthful. However, interestingly, this reduction in compensation does not negatively affect the efficiency of the relationships. When sellers move first within the transaction, they accept the altered terms under incomplete information and are willing to provide similar quality as under complete information, despite the reduction in compensation. Consequently, Step (S4) breaks down: Sellers accept lower prices for a given quality in the incomplete information treatment. They do not react by lowering quality, as they would under complete information.

Our paper contributes to the understanding of the relevance of relational knowledge for building and maintaining relational contracts. When sellers have an informational advantage and move second in the relationship, the relation unfolds as if the cost signal was hard information. If, however, buyers move second, the terms of the relational contract rather seem to reflect the buyers' expectation about the sellers' payoffs from trade and thereby his doubts about the truthfulness of the cost signal. In both cases, however, the efficiency of the relationship is unaltered. Our findings therefore imply that reduced relational knowledge from informational incompleteness can affect relational contracts, although, at least in our setting and with our manipulation through incomplete information, not in a way that harms efficiency.

To our knowledge, our paper is the first to provide empirical evidence on the formation of relational contracts with exogenous variation in asymmetric information. In other words, the paper investigates the interplay of credibility and clarity in forming relational contracts. More generally, the empirical knowledge about the determinants of successful relational contracts is still small. Brown et al. (2004, 2012) study relationship formation in markets under complete information but with varying market power. Brown et al. (2004) show that bilateral long-term relationships in which rents are shared emerge endogenously in markets with excess supply of labor. In Brown et al. (2012), they show that these results extend to markets with excess demand for labor, but bilateral long-term relationships are less frequent. Camerer and Linardi (2012) extend the setup of Brown et al. (2004) and include stochastic hiring shocks in which firms cannot hire, and they show that relationships are robust to such shocks. Renner and Tyran (2004) study buyer-seller relationships in markets for experience goods when temporary cost shocks change the terms of the implicit contract. They show that beneficial long-term relationships prevail in such settings, but are prone to price stickiness. Our paper differs from this literature by dealing with persistent asymmetric information about cost types in relational contracts.

The evidence presented in our paper also informs a growing theoretical literature

analyzing repeated buyer-seller relationships with persistent hidden information, usually specified as principal-agent relations (Levin, 2003; Halac, 2012; Yang, 2013; Li and Matouschek, 2014; Malcomson, 2015). However, there are several differences between this literature and our set-up which aims at testing the pure effect of reduced relational knowledge on the terms of relational contracts. The models studied in this literature focus on the following questions: Under which circumstances can cooperative equilibria with high quality levels be sustained? How do the equilibrium dynamics look like? To which extent are the equilibria separating, that is, reveal the private information to the uninformed party? Most importantly, contrary to our paper, the existing literature focuses mainly on cases where the efficient actions depend on types, so that efficiency requires separation.² However, there is an important common element: Some types of the privately informed parties have an incentive to misreport. For instance, in Levin (2003), the principal wants to understate profits; in Malcomson (2015), he wants to overstate costs; in Halac (2012) he may have incentives to over- or understate outside options. Thus, the potential for mistrust, which we emphasize, is also present in these models.

The remainder of the paper proceeds as follows: Section 2 describes the details of our experimental implementation. Section 3 derives hypotheses. Section 4 presents our empirical results for trust contracts. In Section 5, we discuss bonus contracts. Section 6 concludes.

2 Experimental Design

Our experimental design modifies Brown et al. (2004) by considering fixed relationships and allowing for persistent asymmetric information. Subjects are randomly allocated to be either sellers or buyers. These roles remain fixed for the whole experiment. A relationship lasts for the whole duration of the experiment, with 15 transactions. In each transaction, the buyer of a good pays a price p and the seller provides this good with

²For non-persistent agent cost types, Levin (2003) shows that stationary contracts are optimal under general circumstances. Moreover, the optimal contracts with asymmetric information do not involve full type separation. Finally, no cost type exerts the first-best effort level. Contrary to Levin, Li and Matouschek (2014) suppose that the principal has (non-persistent) private information about the profitability of the firm. They show that bad states lead to conflicts that build up gradually, but are resolved immediately in good states. In a setting with persistent agent cost types drawn from a continuum, Malcomson (2015) shows that there are no fully separating perfect Bayesian equilibria satisfying standard refinements – this differs from our setting with discrete types. Halac (2012) focuses mainly on private information about the (constant) outside option of the principal. She shows that separating equilibria require a high prior probability of uncommitted types (with low outside options). She also discusses equilibrium dynamics, in particular, the speed of information revelation. Like Malcomson (2015), Yang (2013) considers agents with private cost information.³ Contrary to the remaining literature, he assumes that principals and agents are rematched after every termination of a relationship. He shows that wages must increase with tenure as fast as expected productivity.

costly quality q . The stage game is structured as a “trust contract”, with the following order of moves: The buyer moves first, paying the price and indicating a desired quality. The seller observes the price and the desired quality, and can then freely choose quality, without having to deliver the desired quality. In other words, quality cannot be enforced.

As our main treatment variable, we vary the buyer’s information about the seller’s cost type. In each experimental session, it is common knowledge that 50 percent of the sellers will be randomly assigned to high and low cost types, respectively. Prior to the interaction, each seller learns his own type. In our complete information treatments, a buyer also learns the seller’s true type before the interaction begins. In the incomplete information treatments, the seller holds private information about his type, and he can send the buyer a cheap talk signal about his type after learning it. The seller can select the signal independent of his true type, and the buyer knows this.

2.1 Parameters

In each transaction, the buyer chooses a price p from $[0, 100]$ and the seller selects a quality q from $\{1, 2, \dots, 9, 10\}$. The desired quality q^d is chosen from the same set. For given choices of p and q within a round, a buyer’s material payoff $\Pi_B(p, q)$ and the seller’s material payoff $\Pi_S(p, q, \theta)$ are given by

$$\Pi_B(p, q) = 10 \cdot q - p \text{ and } \Pi_S(p, q, \theta) = p - c(q, \theta)$$

where $c(q, \theta)$ indicates the cost of quality q given cost type $\theta \in \{L, H\}$ as summarized in Table 1 below. Costs are strictly increasing and marginal costs are weakly increasing in quality, irrespective of the seller’s cost type. For any $q > 1$, $c(q, L)$ is strictly lower than $c(q, H)$. The difference between high and low costs, $c(q, H) - c(q, L)$, is increasing in q . Yet, since the marginal benefit of quality for a buyer always strictly exceeds the marginal cost of quality for her seller, the efficient quality is 10 under both cost regimes.

Table 1: Cost of Quality

q	1	2	3	4	5	6	7	8	9	10
$c(q, L)$	0	0.5	1	2	4	6	8	10	13	16
$c(q, H)$	0	3	6	10	15	20	25	30	36	42

In all experimental conditions, and before the interaction took place, every seller is randomly assigned to either type L or type H and then privately informed about it. After observing their type, sellers within the incomplete information conditions choose between the message “I have low costs.” and the message “I have high costs.”, irrespective of their actual type. Hence, sellers are free to either be honest or to lie about their

type. A message can only be chosen at the beginning of the experiment and cannot be reversed later. A buyer receives the message selected by her seller and is informed that she will never obtain definite information about her seller’s true costs. In the complete information conditions, a buyer receives either the message “Your seller has high costs.” or the message “Your seller has low costs.”, depending on the true costs of her seller.

2.2 Procedures

Assignment to the roles of a buyer and seller is random, the matching of buyers and sellers is randomly determined, and each match persists for fifteen rounds. At the end of each period, both buyer and seller receive a summary of their choices in the current round including the price and quality as well as the desired quality. Every subject is additionally informed about the own material payoff in the current round in terms of the experimental currency “Punkte” (points). The sum of payoffs, taken over all rounds, is converted into real money at the end of the experiment (10 points=1 CHF(\$1)) and paid out with the show up fee (10 CHF). In the incomplete information experiments, we additionally elicited buyers’ first-order beliefs about the accuracy of the cost signal after the last interaction. We furthermore elicited sellers’ second order beliefs about their buyers’ first order beliefs.⁴

Table 2: Overview: Treatments, Sessions and Participants

Treatment	Number of Sessions	Total Number of Subjects
Trust Complete Information (TC)	4	116
Trust Incomplete Information (TI)	4	128

The experiment was computerized using the software z-tree (Fischbacher, 2007). For organizing and recruitment, we used the software hroot (Bock et al., 2014). Our subject pool consists primarily of students at the University of Zurich and the Swiss Federal Institute of Technology in Zurich. In total, 244 subjects participated in the experiment between Fall 2013 and Summer 2014. No subject participated in more than one session. On average, a session lasted 95 minutes with an average payment of 44.7 CHF (\$ 45). An overview of the treatments and number of subjects and sessions is shown in Table 2.

⁴We asked each buyer about all sellers’ message choices, rather than about his actual seller. For instance, in a session with 16 sellers where 8 are assigned to low costs, we asked the following question: “8 out of 16 sellers were assigned to low costs. How many of these sellers with true low costs sent the message “I have low costs.” to their buyers?”. In a session with 16 sellers where 8 are assigned to low costs, sellers answered the following question: “Your buyer was asked the following question: ‘8 out of 16 sellers were assigned to low costs. How many of these sellers with true low costs sent the message “I have low costs.” to their buyers?’ What do you believe: which answer did your buyer provide in response to this question?”. Elicitation was incentivized. Subjects earned an extra 20 points for each question if their stated belief was correct.

3 Hypotheses

Our finite horizon game with complete information clearly has a unique subgame-perfect equilibrium with minimal qualities and prices in every period; moreover, there are no perfect Bayesian equilibria with cooperation in the incomplete information game. Previous repeated games experiments suggest, however, that cooperation is likely to arise. There are at least two approaches that would predict such cooperation.⁵

First (see Appendix, Section A), admittedly involving a substantial leap of faith, one can take an as-if approach by arguing that subjects have limited foresight and hence behave in a way that is broadly consistent with equilibrium behavior in infinitely repeated versions of the game, except for end game effects.⁶ The advantage of this approach is that it generates simple results on the equilibrium sets. For instance, it is straightforward to characterize the set of quality-price combinations that are sustainable as trigger-strategy equilibria for complete information for any given discount factor. In an incomplete information version of the game, one can consider the set of pooling equilibria in trigger strategies for which both high cost and low cost types choose the same qualities. Interestingly, the prices and qualities that can be sustained in such equilibria correspond exactly to those that emerge as trigger strategy equilibria of the complete information game with high costs. Credibility is therefore unaffected and, in this sense, the infinite horizon approach provides no reason why behavior in high cost relationships with complete and incomplete information should differ.

Alternatively (see Appendix, Section B), one can take the finiteness of the time horizon seriously, but instead argue that some players are committed to fair behavior.⁷ We suppose that players are privately informed about whether they have such a “fair type” rather than a standard “selfish type”. Specifically, we assume that a fair type chooses an equal split of the surplus as long as he is not certain that the other player is a selfish type.⁸ Thus, even the case with complete cost information becomes a game of incomplete information. For sufficiently high shares of fair types, there exists an equilibrium with pooling at the highest quality level where the selfish types imitate the fair types until the penultimate period, and, in the last period, the selfish buyers still

⁵For technical details, we refer the reader to the Appendix.

⁶See Mengel (2014) for one formalization of limited forward-looking behavior.

⁷In other contexts, commitment types have been used for a long time ((Kreps et al., 1982))

⁸Thus, fair sellers choose qualities such that the surplus calculated on the basis of this quality and the previous price of the buyer is the same for both players. For instance, the low type who has received a price of 53 chooses the efficient action 10, which generates a buyer surplus $100-10-53=37$ and a seller surplus $53-16=37$. Fair buyers choose prices such that the surplus calculated on the basis of this price and the previous quality of the buyer is the same for both players. For instance, a fair buyer who observes quality 10 in the previous period sets a price of 53. In period 1, fair buyers choose a price that would split the surplus in half for the highest quality choice of the seller, that is, they give the seller the benefit of the doubt.

pool with the fair buyers so as to induce high qualities from the fair sellers.

What happens in this game when incomplete cost information is introduced? In this case, the buyer can no longer know with certainty which price is “fair” in the sense that it splits the surplus equally. We consider two extreme cases in which the buyer and seller mutually agree on a fair price.

First, suppose that the high-cost seller considers the complete information fairness price (for high costs) as adequate response to a high cost signal. Suppose further that the fair buyer is prepared to pay this fairness price. We thus think of the buyer as being credulous, accepting the signal at face value. In this case, the game has an efficient pooling equilibrium under exactly the same conditions as with complete information. Intuitively, as the buyer is prepared to accept the high signal as if it were verifiable, everything is as under complete information.⁹ This suggests that incomplete information might have no effect on efficiency.

Second, assume that the buyer is skeptical, that is, she thinks the signal is completely uninformative about the cost type (Appendix B.2.1). As a result, she is prepared to pay a fairness price that equalizes expected surplus, i.e., a weighted average of the fairness price under complete information with high and low costs, respectively. Suppose that the seller has the same perception of what a fair price is. In such a setting, it turns out that the efficient pooling equilibrium is easier to sustain than under complete information: The required threshold share of fair sellers is lower, because the high-cost sellers understand they cannot ask for as much as under complete information. Consequently, it might even become easier to sustain efficient pooling equilibria with high cost types.

However, with incomplete cost information, disagreements about what constitutes the fair price can arise easily, certainly more easily than in the case in which the generated surplus is common information. If buyer and seller do not agree on the fair price, sustaining the relational contract becomes difficult. To illustrate this scenario, consider a setting where fair sellers are committed to surplus sharing and both players regard a weighted average of the complete information fairness prices for high and low costs as adequate. However, they disagree about weights. A seller who knows he has high costs regards a higher price as fair than a buyer is prepared to pay when he is uncertain about

⁹In this argument, we have ignored the possibility that fair sellers also have a preference to truthfully reveal their cost type. If this were the case, unfair low cost sellers could be more likely to send a high cost signal. In consequence, the share of fair sellers conditional on a high cost signal might be reduced. However, it is not clear at all whether fairness and honesty are correlated in such a setting. It has been shown that people exploit so called “moral wiggle room” (Dana et al., 2007; Lazear et al., 2012) in order to behave more selfishly. Applied to our game, fair low cost sellers exploit the lack of transparency with respect to costs in the cost signal stage to gain higher rents in the relationship. Once they engage in the buyer-seller relationship, however, there is transparency with respect to fairness conditional on the signal, so fairness preferences of sellers induce them to behave like high cost fairness commitment types.

the true costs. Even small disagreements of this kind can lead to an equilibrium where cooperation slowly breaks down. After a high cost signal, the buyer pays a slightly lower price than a high-cost seller considers adequate for the highest quality level. As the resulting quality of the seller does not quite satisfy the buyer, she responds with a slight price reduction. By iteration, a downward spiral of qualities and prices emerges in equilibrium.

In general terms, the analysis shows that, as long as trading parties can agree on the fair price, the constraints on sustaining a relational contract following a high cost signal are either the same or weaker than under incomplete information, compared to the complete information high cost case. However, incomplete information introduces ambiguity about what price is fair. The implied potential for a mismatch in fairness views between buyers and sellers can then lead to a breakdown of cooperation, as outlined in our steps (S1) - (S4) in the introduction.

Summing up, our theoretical analysis shows that credibility in high cost relationships remains unaffected, and hence there is no reason to believe that high cost relationships become harder to sustain *due to reduced credibility*. However, ambiguity about the generated surplus can make it harder to coordinate on a fair price for a given quality, which can cause cooperation breakdowns. Consequently, high cost relationships may become harder to sustain *due to reduced clarity*.

We can summarize our Null hypothesis, or *credibility hypothesis*, as follows:

Hypothesis 1. *The average quality in actual high cost relationships is the same with incomplete information and with complete information.*

Alternatively, we have argued that a lack of clarity may result in a reduction in efficiency in relational contracts, implying the *clarity hypothesis*:

Hypothesis 2. *Incomplete information lowers the average quality in high cost relationships relative to complete information.*

While the effect of incomplete information is our main focus, we also present some additional hypotheses which are natural by-products of the analysis.

Hypothesis 3. *There is a positive relation between price and quality for complete as well as incomplete information.*

This hypothesis can be derived from the repeated games approach as well as from fairness-based models. In the former setting, incentive constraints imply that the lowest price that is consistent with equilibrium behavior is increasing in quality, reflecting the higher costs for the seller. In the latter setting, surplus splitting directly implies that high prices induce high qualities, and high qualities induce high prices. The next hypothesis deals with cost effects:

Hypothesis 4. (i) *Quality is higher with low costs than with high costs.*

(ii) *An increase in quality corresponds to a greater increase in price in high cost relationships than in low cost relationships. Also, prices for any given quality level are higher with high costs than with low costs.*

Intuitively, (i) with higher costs, for any given quality level the seller's incentive constraints for an equilibrium in the infinitely repeated game can only be fulfilled with a higher price when costs are high. This makes it hard to achieve high qualities. Fairness-minded players will regard lower qualities as adequate for any given price level when costs are higher. Similarly, (ii) can be based on the standard repeated games logic as well as on fairness considerations.

4 Results

We have argued that incomplete information might lead to disagreements about fair prices and thus makes it harder to sustain relational contracts in high cost relationships, despite the fact that credibility constraints are not tightened. In this section, we present our main results and relate them to the four steps (S1) - (S4) that we outlined in the introduction.

The two panels in Figure 1 show average quality over time for high and low cost relationships, conditional on whether costs are publicly or privately known.¹⁰ Because our hypotheses are derived for high cost relationships, we will first focus on these. Later we will discuss some insights for low-cost relationships. The left panels strongly suggests that incomplete information has no systematic effect on quality in high cost relationships. Average quality on high cost relationships under complete information is 5.4, and 5.6 under incomplete information. This difference is insignificant, using a t-test (p-value: 0.75).¹¹ This observation leads to our first result:

Result 1 (The Effect of Incomplete Information). *Incomplete information has no significant impact on average quality in high cost relationships.*

Our data therefore does not allow us to reject our *credibility hypothesis*. Contrary to the clarity hypothesis, reduced relational knowledge through incomplete information seems to have no effect on the average quality in high-cost relationships. To understand where our four-step logic failed, we will consider each step in detail.

¹⁰Here, we ignore the cost signal sent by the seller at the beginning of the relationship. However, high cost sellers are our main focus, and all but one high cost seller indicated high costs.

¹¹The t-test is conducted using the average quality provided in a trading relationship as one observation. Hence, there are 29 observations from complete information high cost relationships and 32 observations from incomplete information high cost relationships.




Figure 1: Effects of incomplete Information

Step (S1) in our chain of reasoning (also captured as Hypothesis 3(ii)) suggests that high cost sellers should receive more compensation for quality than their low cost counterparts under complete information. Our data confirms this intuition:

Result 2 (Price-Quality Relationship). *(i) There is a positive relation between average price and average quality for any cost level. (ii) Quality is significantly less sensitive to prices in low cost relationships than in high cost relationships.*

Evidence for this result is shown in Figure 2, which plots average prices and average qualities for every trading relationship conditional on costs. The quality-price relationship is positive for both cost structures, and it appears to be flatter for low cost relationships. Column (1) of Table 5 confirms this observation. In a Tobit regression of paid prices in the complete information treatment, we first see a large and significant coefficient of 7.4 on *Quality*, indicating that one unit of higher quality is compensated with a 7.4 points higher price in high cost relationships. Moreover, the interaction between quality on low costs is -1 , implying that the slope of the price-quality relationship is flatter under low costs. The difference is significant ($p < 0.05$).

We now consider step (S2), which argues that buyers who observe a high cost signal are not sure whether they indeed face a high cost seller or a low cost seller who is trying to gain an unfair advantage.

Table 3 shows the relative frequency of high cost signals, conditional on true costs. High cost sellers almost unanimously indicate high costs, as do roughly two thirds of the low cost sellers. Further, the table shows that the buyers approximately expected this frequency, as indicated by the first order beliefs, and sellers believed that buyers believed that low cost sellers would indicate high costs with roughly the actual empirical




Figure 2: Association between price and quality. Each dot represents one relationship and depicts the average price and quality provided over the 15 periods in the relationship.

frequency (as indicated by the second order beliefs).¹² Hence, the empirical frequency of high cost signals by low cost sellers coincides with the beliefs of buyers as well as sellers.¹³

Table 3: Fraction of High Cost Signals under Incomplete Information

	Trust Contract
Overall	81.3 %
High Cost Seller	96.9 %
Low Cost Seller	65.6 %
1 st order beliefs (low cost sellers)	71 %
2 nd order beliefs (low cost sellers)	63 %

Consistent with step (S2), the belief data therefore implies that buyers had (justified) doubts about the cost type of a seller after observing a high cost signal. When buyers

¹²t-tests show that the actual frequencies of signaling high costs are not significantly different from the elicited first order beliefs of buyers ($p = 0.53$, t-test) and the second order beliefs of sellers ($p = 0.76$, t-test).

¹³For the more obvious case of high cost sellers, it was generally expected that they truthfully signal high costs, which is also what is empirically observed.

are confronted with a high cost signal in the first period, they also translate these doubts into action by asking for higher quality increases in response to any price increase than buyers in the complete information treatment who are matched with high cost sellers. This is evident in column (1) of table 4, which shows results from a Tobit regression of desired quality on a complete information treatment dummy, interacted with the price paid up front.¹⁴ In column (1), we focus on the first period. As there has been no previous interaction, but the buyer already received the cost signal, the first period provides the cleanest indication for the quality that buyers consider appropriate for the price paid by them. Here, we include all buyers who observed a high cost signal in the regression, independent of their sellers' true costs. From the complete information treatment, only high cost relationships are included.

Table 4: Desired quality levels conditional on paid prices and information treatment

	(1)	(2)
Complete information treatment	2.082 (1.607)	-0.710 (1.398)
Paid price	0.127*** (0.016)	0.126*** (0.013)
Complete information*price	-0.064* (0.037)	-0.008 (0.028)
Constant	2.558*** (0.684)	3.950*** (0.666)
Pseudo R^2	0.165	0.155
Observations	81	1215

Tobit Regressions on desired quality on observations for which high costs were common knowledge (complete info treatment) or high costs were signaled (incomplete info treatment). Lower limit: 1; upper limit: 10. Column (1) only contains observations from the first period. Column (2) contains observations from all periods. In column (2), standard errors are clustered at the relationship level (81 clusters)

Column (1) shows that, for 10 additional points paid up front to the seller, buyers demand .64 units of quality less when they know for sure that the seller has high costs, compared to the case in which they received the non-verifiable high cost signal. This interaction is significant at the 10% level. However, it turns out that, contrary to our conjecture in Step (S3), these doubts did not translate into lower prices for a given quality.

Column (2) of Table 5 shows tobit regressions on prices that control for the different treatments, taking into account cost signals and quality. CL is a dummy for low cost relationships under complete information. IH denotes a dummy for high cost relation-

¹⁴We use Tobit regressions for all our regression analyses of chosen prices and quality levels to account for corner solutions, which appear frequently, in particular for quality choices.

Table 5: Prices paid by treatment and cost signal

	(1)		(2)	
complete info - low cost (CL)	-1.372 (3.940)		-1.081 (4.043)	
Quality	7.420 (0.415)	***	7.458 (0.423)	***
CL X Quality	-1.000 (0.484)	**	-1.043 (0.499)	**
incomplete info - high cost (IH)			0.939 (4.073)	
incomp. info - low cost - low cost signal (ILL)			-4.021 (4.092)	
incomp. info - low cost - high cost signal (ILH)			0.480 (4.524)	
IH X Quality			-0.197 (0.518)	
ILL X Quality			-1.365 (0.552)	**
ILH X Quality			0.091 (0.553)	
Constant	1.429 (3.000)		10.573 (0.593)	***
Period Fixed Effects?	Yes		Yes	
Pseudo R^2	0.202		0.195	
Observations	870		1830	

Tobit Regressions on paid prices. Lower limit: 0; upper limit: 100. To control for end game effects, period fixed effects are included in the regression. Standard Errors are clustered at the relationship level (122 clusters). Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

ships under incomplete information, independent of the cost signal.¹⁵ *ILL* is a dummy for low cost sellers who sent a low cost signal, and *ILH* is a dummy for low cost sellers who signaled high cost in the incomplete information treatments. The baseline category consists of high cost relationships under complete information.

The highly significant coefficient on quality again captures the positive relation between quality and prices in high cost relationships under complete information. The interactions of our treatment and signal dummies with quality are of particular interest, since they directly indicate whether the quality-price relationship is different from the one in the baseline. Compared to the baseline treatment with high costs and complete information, the quality-price relationship is significantly flatter in low cost complete information relationships and in relationships in the incomplete information treatment

¹⁵This signal was high in all but one case.

in which low cost sellers indicated low costs. However, such a difference to the baseline neither arises for high cost relationships under incomplete information nor for low cost relationships under incomplete information in which the seller indicated high costs. This implies that the compensation of sellers who signal high costs under incomplete information is similar as for sellers with high costs under complete information. Moreover, the coefficient on the interaction for sellers who signal low costs under incomplete information is not significantly different from the coefficient on the interaction for actual low cost sellers under complete information. Thus, we obtain the following result.

Result 3. [*Quality-Price Relationships under Incomplete Information*] *In the trust contract games with incomplete information, buyers compensate sellers as if the cost signal was truthful.*

The result shows why incomplete information has no adverse effect on efficiency: Even though buyers doubt that high costs are truthful, they behave as if they were taking them at face value. This also becomes evident in column (2) of table 4, where observations from all periods are considered. The sizeable and significant negative interaction between complete information and paid prices from period 1 essentially disappears when all periods are considered. Consequently, buyers appear to accept similar terms than in the complete information case. This behavior helps to maintain the same efficiency levels as under complete information.

Our data from low cost relationships also provides some noteworthy insights. First, Figure 1 suggests that the average quality is lower in high cost relationships than in low cost relationships under complete information. The difference in average quality between low and high cost relationships with complete information amounts to 1.35 points ($p = 0.03$, Mann-Whitney test).¹⁶ Consistent with Hypothesis 3(i), the lower prospective future rents in high cost relationships have a detrimental effect on cooperation and efficiency.

Moreover, at first sight Result 3 seems to suggest that low cost sellers profit from signaling high costs under incomplete information. After all, it enables them to obtain higher prices for a given quality. However, this conjecture is not true, because costs affect quality as well as prices: Low cost relationships with high cost signals turn out to be less efficient than low cost relationships with low cost signals.

Table 6 provides evidence for this finding. It shows results from a Tobit regression of quality on dummy variables for different combinations of true costs and cost signals.¹⁷

¹⁶Average quality is higher in low cost relationships in 14 out of 15 periods. The quality difference between low and high cost relationships is also confirmed in a tobit regression that accounts for the frequent corner observations at quality levels of 1 and 10 as well as for period fixed effects that account for the end game effect. The coefficient on the high cost relationship dummy in this regression is significant at the

Table 6: Tobit Regressions: Incorporating the Cost Signal

	(1)	
complete info - low cost (CL)	2.073	*
	(1.108)	
incomplete info - high cost (IH)	0.232	
	(1.052)	
incomp. info - low cost - low cost signal (ILL)	3.371	**
	(1.647)	
incomp. info - low cost - high cost signal (ILH)	0.492	
	(1.161)	
Constant	5.208	***
	(0.812)	
Adj. R^2		
Observations	1830	

Tobit Regression on provided quality. Lower limit: 1; upper limit: 10. Standard Errors in parentheses, clustered at the relationship level (122 clusters). Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Relationships in which the seller sent a low cost signal under incomplete information are significantly more efficient than high cost relationships under complete and incomplete information ($p = 0.05$ for the latter). Low cost relationships in which a high cost signal was sent, however, are not significantly more efficient than actual high cost relationships, and they are also significantly less efficient than low cost relationships in which a low cost signal was sent ($p = 0.08$). There are two potential explanations for these patterns in the low-cost relationship data. First, reflecting our discussion of credulous buyers in Section 4, it may be that behavior and outcomes in relationships under incomplete information are as if the cost signal is truthful, which would explain why relationships under incomplete information conditional on the cost signal so closely mirror their complete information counterparts. Second, it may be that the cost signal has no efficiency effect, but that the efficiency difference is the consequence of selection of specific seller types with different abilities to maintain efficient relationships based on the cost signal.¹⁸

Our data also allows us to assess distributional consequences of incomplete informa-

$p = 0.06$ level.

¹⁷The analysis of the impact of cost signals should not be understood as causal, since cost signals are endogenous to individual characteristics, and hence selection may take place conditional on the signal. Nevertheless, the associations in the data are interesting.

¹⁸These hypotheses could be analyzed in future experiments. For example, one could conduct a within subjects design in which subjects participate in both a complete and the incomplete information treatment. This would allow to test the hypothesis that those sellers who are involved in less efficient relationships under complete information are more likely to signal high costs under incomplete information. Another possibility would be to conduct a treatment in which the cost signal is random. If there is a treatment effect related to the high cost signal in the latter case, selection cannot be the explanation. Since these questions are not the main focus of this paper, we leave these suggestions for future research.

tion. While we have already seen that efficiency in high cost relationships is unaffected, it will turn out that the terms of the relational contract have changed and that buyers and sellers coordinate on different equilibria under incomplete information resulting in differences in distributional outcomes. From Figure 2 above, we see that prices are generally higher than the equal split price, indicated by the solid lines. Thus, sellers get a disproportionately large share of the surplus: Table 7 summarizes the average profits earned by buyers and sellers across the different treatments.

Table 7 powerfully confirms the pattern already visible in Figure 2. Sellers capture significantly larger rents than buyers, except in low cost relationships under incomplete information in which low costs were signaled. The largest payoff difference materializes in low cost relationships under incomplete information in which high costs were signaled. However, somewhat stunningly, the table also reveals that low cost sellers are in fact not better off signaling high rather than low cost. The fact that relationships on average feature a lower quality following a high cost signal fully makes up for the higher prices conditional on quality, so that sellers payoffs under complete and incomplete information are about the same. The efficiency loss relative to complete information low cost relationships is fully born by the buyers, who make significantly smaller profits when the seller signals high rather than low costs.¹⁹

Table 7: Average Profits by Treatment and Cost Signal

	Trust Contracts	
	Buyer	Seller
Complete Info - High Cost	14.27	20.88
Complete Info - Low Cost	26.44	32.24
Incomplete Info - High Cost	15.27	20.98
Incomplete Info - Low Cost	22.60	32.48
Incomplete Info - Low Cost - Low Cost Signal	29.32	31.93
Incomplete Info - Low Cost - High Cost Signal	19.08	31.34

One potential reason for the high payoff share of the sellers might be that they move second and therefore control the distribution of rents. This suggests that the order of moves may play an important role in selecting the terms of the contract. We now deal with this point in detail.

¹⁹ Again, this is an association and not a causal relationship. It could be the consequence of selection, as discussed before. Considering low cost relationships independent of the cost signal, buyers on average earn 4 points less per period under incomplete information than under complete information, but this difference is not statistically significant ($p = 0.15$).

5 Bonus Contract Games

Since our hypothesis on possible adverse effects of reduced relational knowledge relates to uncertainty about rent distribution, the uninformed party may require power over the distribution of rents within a transaction to translate its uncertainty into meaningful action. To test the role of second-mover rent control, we conducted a second experiment in which the buyer moves second and therefore directly controls the distribution of rents through his action.

5.1 Experimental Design: The Bonus Contract Game

We carried out eight sessions of a “bonus contract game”. The bonus contract game is similar to the trust contract game, except that the order of moves within the stage game is reversed. In every period, the seller moves first, chooses quality and incurs the associated cost. In addition, he indicates a desired price. The buyer observes the quality and the desired price, and then he chooses the price he wants to pay. Thus, again neither quality nor price is contractible.

As before, the stage game is repeated for 15 periods. We conducted 4 sessions with complete information and 4 sessions with incomplete information. As in the trust contract games, information concerns the cost type of the seller, which was either “high” or “low”. In the complete information treatments, the seller’s cost type was common knowledge, whereas in the incomplete information treatments, the seller could send a non-verifiable cost signal.²⁰ In total, 252 subjects participated in the additional experiments in 2014 and 2015. No subject participated in more than one session, and no subject had previously participated in the trust contract treatments. On average, a session lasted 95 minutes with an average payment of 52 CHF (\$ 52).

5.1.1 Hypotheses

As argued above, there are reasons to believe that the effect of the information structure may differ for bonus and trust contracts. Specifically, it appears plausible that the effect of a lack of payoff clarity crucially depends on whether the buyer has the necessary control over the rent distribution to translate his doubts into meaningful action. Such control requires that the buyer only pays the price after receiving the product, as in the bonus game. Thus, we expect that, contrary to the trust contract games, Step (S3) in the reasoning outlined in the introduction may hold for bonus treatments and that this sparks off the sellers’ reaction described in Step (S4).

²⁰ Again, the experiment was computerized using the software z-tree (Fischbacher, 2007). For organizing and recruitment, we used the software hroot (Bock et al., 2014).

5.2 Results

Figure 3 shows average quality in low and high cost bonus contract relationships during the 15 periods, both under complete and under incomplete information. Again, to keep credibility constant, we assess our main hypothesis by only considering high cost relationships. Average quality in complete and incomplete information treatments is similar, 7.39 and 7.07, respectively. The difference is not statistically significant ($p = 0.51$, Mann-Whitney Test). Therefore, as in the trust contract experiments, the lack of payoff clarity does not appear to have a negative impact on efficiency in relational contracts.

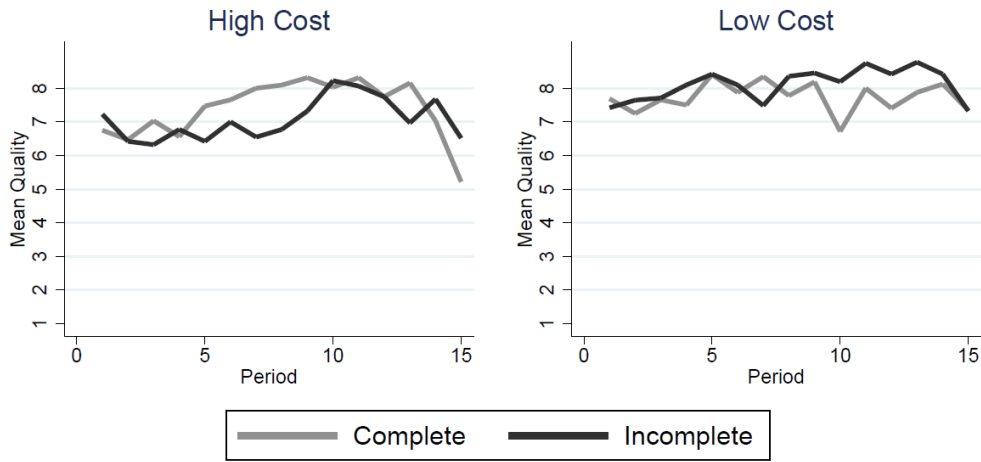


Figure 3: Quality over time in Bonus Contract relationships. Left Panel: High Cost relationships under complete and incomplete information. Right Panel: Low Cost relationships under complete and incomplete information.

To assess differences between the trust and bonus contract experiments, we turn to regression analysis. Table 8 analyzes the determinants of quality using Tobit regressions, containing observations from both the trust contract and the bonus contract experiments. Column (1) shows coefficients of a Tobit regression of quality in high cost relationships on treatment dummies and their interactions. If, as hypothesized, the second mover advantage affects quality, the bonus treatment dummy and the incomplete information treatment dummy should interact negatively. Column (1) confirms that this interaction is indeed negative and roughly equal to 0.9 quality points. This suggests that moving from the trust contract games to the bonus contract games, an efficiency-reducing effect of incomplete information may be present, but the interaction is not significant. Consequently, we cannot reject the hypothesis that incomplete information has no effect on efficiency, no matter who controls rent distribution in the stage game. A lack of payoff clarity therefore does not seem to negatively impact the efficiency of

relational contracts.²¹

Result 4. *[Incomplete Information and Bargaining Power] Keeping credibility constraints constant, incomplete information has no significant effect on the efficiency of relational contracts, independent of who controls the rent distribution in the stage game.*

Table 8: Tobit Regressions: Effect of Incomplete Information on Quality

	(1)		(2)	
	High Cost		Low Cost	
Incomplete Info	0.234 (1.109)		−0.652 (1.093)	
Bonus Contract	3.808 *** (1.199)		2.529 ** (1.173)	
Bonus X Inc. Info	−0.891 (1.596)		0.954 (1.537)	
Constant	5.178 *** (0.858)		7.313 *** (0.772)	
Pseudo R^2				
Observations	1860		1860	

Tobit Regressions on provided quality. Lower limit: 1; upper limit: 10. Standard Errors are clustered at the relationship level. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 8 provides a further interesting insight that is independent of our incomplete information manipulation but clearly noteworthy: The coefficient on the bonus contract dummy in Columns (1) and (2) in Table 8 reveals that bonus contract relationships perform considerably and significantly better than trust contract relationships.

Turning back to the impact of incomplete information on efficiency, it is again informative to analyze our four steps in detail and to assess which of these do not apply in the bonus contract game. To this end, figure 4 shows a scatter plot of the average quality and price in each complete information bonus contract relationship over the course of the 15 periods. The solid lines again depict the price that would lead to an equal split of the rent, given average quality.

As in the trust contract case, higher quality is rewarded with higher prices, and sellers with higher costs are rewarded with higher prices for a given quality than low cost sellers. This can be seen in Table 9, which shows results from a tobit regression on paid prices. The coefficient on *Quality* shows the significant and positive relationship between average prices and quality in high cost relationships under complete information. The coefficient on the interaction between complete information low costs and qual-

²¹Column (2) in Table 8 shows that incomplete information does not have a significant impact on average quality in low cost relationships either, even though the cost signals are not all truthful.




Figure 4: Association between price and quality in bonus contract relationships. Each dot represents one relationship and depicts the average price and quality provided over the 15 periods in the relationship.

ity is significantly negative, indicating that the quality-price relationship is significantly flatter for low cost relationships under complete information, consistent with (S1).

Moreover, similar to the trust contract games, all high cost sellers and 77.4% of the low cost sellers signal high costs in the bonus contract games. First and second order beliefs are again well aligned with these numbers.²² Thus, buyers have significant (and justified) doubts about true costs when they observe a high cost signal, consistent with (S2).

Further, we argued that buyers might change their behavior in the bonus contract game because they control the distribution of rents in the stage game. First, as conjectured, Figure 4 shows that the second mover controls the rent distribution in the stage game: Quite strikingly, almost all observations are now to the right of the equal rent split line. This implies that, on average, the seller receives a lower price than the rent splitting price, and the buyers consequently receive a larger share of the rents.

²²Buyers on average believe that 69% of low cost sellers would signal high costs. Sellers second order belief about this belief is 68%. These beliefs are not significantly different from the actual frequency ($p = 0.28$ and $p = 0.23$, respectively; t-tests).

Table 9: Prices paid in Bonus Contract Relationships

	(1)	
complete info - high cost (CL)	0.297	
	(2.203)	
incomplete info - high cost (IH)	0.171	
	(2.333)	
incomp. info - low cost - low cost signal (ILL)	1.329	
	(2.308)	
incomp. info - low cost - low cost signal (ILH)	3.941	
	(4.504)	
Quality	7.467	***
	(0.161)	
CL X Quality	-1.498	***
	(0.320)	
IH X Quality	-0.686	**
	(0.322)	
ILL X Quality	-1.678	***
	(0.406)	
ILH X Quality	-1.302	***
	(0.346)	
Constant	-8.186	***
	(1.410)	
Adj. R^2	0.165	
Observations	1890	

Tobit Regressions on paid prices. Lower limit: 0; upper limit: 100. To control for end game effects, period fixed effects are included in the regression. Standard Errors are clustered at the relationship level (126 clusters). Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

But do they use this additional control to translate their doubts in the incomplete information treatment into lower prices? In Table 9, we see that the quality-price relationship changes significantly when information is incomplete. In actual high cost relationships, buyers on average receive .69 points less compensation for a one point quality increase than under complete information, and this difference is significant.²³ Therefore, incomplete information has a clear effect on the quality-price relationship in the bonus contract games. High cost sellers have to accept a lower compensation for the same quality under incomplete information than under complete information. The buyers' doubts about the truthful revelation of high costs translate into lower prices for

²³Moreover, the slope of the quality-price relationship is 1.3 points smaller in relationships with low cost sellers that signalled high costs, and there is no difference in the quality-price relationship for low cost sellers under incomplete information conditional on the signal. Finally, for low cost relationships with high cost signals, the quality-price relationship is also not significantly different ($p = 0.14$) from the one of high cost relationships under incomplete information.

a given quality.²⁴

Result 5. *[Quality-Price Relationships under Incomplete Information] In bonus contract games, sellers indicating high costs under incomplete information get paid significantly less for the same quality than under complete information.*

All told, (S1)-(S3) hold: When buyers know that costs are low, they provide less compensation for quality than when they are high. This gives low cost sellers an incentive to mimic high-cost buyers, which leads to doubts of the buyers regarding the truthfulness of high cost signals. As expected, they are therefore less willing to compensate high-cost sellers for quality increases.

The fact that these lower prices did not result in lower average quality means that the final step (S4) in our reasoning is not borne out empirically: Sellers accepted the lower prices without in turn lowering their quality provision. Relationships coordinated on a different quality-price pair than under complete information. It appears that sellers understood the doubts of the buyers and their price reaction, and therefore did not react to the lower prices by reducing quality. Consequently, efficiency was unaffected.

The differences in control over rent distribution translate into considerable profit differences. Table 10 summarizes profits in the bonus contract games. Incomplete information does not harm buyers in high cost relationships. To the contrary, their profits are significantly higher with incomplete information. This reflects lower prices for a given quality in the incomplete information treatment, and the fact that these lower prices did not lead to reduced quality. The increase in buyer profits therefore corresponds almost one-to-one to a reduction in seller profits, whose profits significantly decrease under incomplete information.

Table 10: Average Profits in the Bonus Contract Games

	Buyer	Seller
Complete Info - High Cost	26.48	18.81
Complete Info - Low Cost	38.48	27.85
Incomplete Info - High Cost	30.12	13.65
Incomplete Info - Low Cost	37.69	31.8
Incomplete Info - Low Cost - Low Cost Signal	40.2	33.53
Incomplete Info - Low Cost - High Cost Signal	36.96	33.39

²⁴In the bonus contract treatments, sellers had the possibility to signal a desired price after they provided the quality up front. In a tobit regression of desired prices on provided quality levels interacted with the information treatment, we find that, in the first period, sellers who signaled high costs desire prices that are even 6.4% larger than the prices desired by high cost sellers under complete information. This difference is, however, not significant. Also when all periods are considered, desired prices are no different between the complete and incomplete information treatments. Hence, sellers would like the buyer to treat the cost signal as truthful. However, as our analysis has shown, they fail to achieve this goal.

Further, we again observe that low cost sellers barely profit from signaling high costs. This result also reflects the lower quality-sensitivity of prices with incomplete information. Since buyers who received a high cost signal are not willing to pay a considerable premium for a given quality, sellers do not profit from signaling incorrect costs.

6 Conclusion

Our experiments have shown that incomplete information reduces clarity about the rents earned by each party in experimental trading relationships. If sellers have private information about their costs, a majority of those with low costs signals high cost, presumably in the expectation of gaining an income advantage. When the stage game is structured as a trust contract game in which the buyer pays a price up front and the seller provides quality subsequently, this lack of clarity does, however, not translate into any differences in behavior. Despite their doubts about the truthfulness of the cost signal, buyers compensate quality as if the cost signal was truthful. Consequently, trading relationships in which high costs are signaled are equivalent in terms of average quality and average prices to full information high cost relationships.

However, if the stage game is structured as a bonus contract so that the buyer moves second, uncertainty about sellers' costs does affect outcomes. Buyers reduce the sensitivity of prices to quality when faced with a high cost signal. They are less willing to compensate sellers for supposedly high costs. But true high cost sellers accept lower prices in the incomplete information treatment. Consequently, the reduced compensation of quality does again not translate into reduced efficiency in the relationship. It does, however, lead to less compensation of higher quality by higher prices and thus to effects on the distribution of rents within the relationship.

The effects of incomplete information on relational contracts are therefore complex. While we find no evidence that the resulting reduction in relational knowledge results in reduced efficiency, buyers and sellers do coordinate on different price-quality combinations conditional on the relational knowledge and the structure of the stage game. We believe that further exploring the determinants of the formation of relational contracts under incomplete information and other sources of lack of clarity is a fruitful avenue for future research.

References

- Bachmann, R. and A. Zaheer (2006). *Handbook of trust research*. Edward Elgar Publishing.
- Bock, O., I. Baetge, and A. Nicklisch (2014). hroot: Hamburg registration and organization online tool. *European Economic Review* 71, 117–120.
- Brandts, J. and G. Charness (2003). Truth or consequences: An experiment. *Management Science* 49(1), 116–130.
- Brown, M., A. Falk, and E. Fehr (2004). Relational contracts and the nature of market interactions. *Econometrica* 72(3), 747–780.
- Brown, M., A. Falk, and E. Fehr (2012). Competition and relational contracts: the role of unemployment as a disciplinary device. *Journal of the European Economic Association* 10(4), 887–907.
- Camerer, C. and S. Linardi (2012). Can relational contracts survive stochastic interruptions? Experimental evidence. University of Pittsburgh, Department of Economics, Working Paper No. 483.
- Dana, J., R. A. Weber, and J. X. Kuang (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* 33(1), 67–80.
- Fischbacher, U. (2007). z-tree - zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2), 171–178.
- Gibbons, R. and R. Henderson (2012). Relational contracts and organizational capabilities. *Organization Science* 23(5), 1350–1364.
- Gibbons, R. and R. Henderson (2013). What do managers do? In R. Gibbons and J. Roberts (Eds.), *Handbook of Organizational Economics*. Princeton University Press.
- Halac, M. (2012). Relational contracts and the value of relationships. *American Economic Review* 102(2), 750–779.
- Kreps, D., P. Milgrom, J. Roberts, and R. Wilson (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory* 27(2), 245–252.
- Lazear, E. P., U. Malmendier, and R. A. Weber (2012). Sorting in experiments with application to social preferences. *American Economic Journal: Applied Economics* 4(1), 136–163.

- Levin, J. (2003). Relational incentive contracts. *American Economic Review* 93(3), 835–857.
- Li, J. and N. Matouschek (2014). Managing conflicts in relational contracts. *American Economic Review* 103(6), 2328–2351.
- Malcomson, J. (2015). Relational incentive contracts with persistent private information. University of Oxford, Department of Economics, Economics Series Working Papers No. 634.
- Mengel, F. (2014). Learning by (limited) forward looking players. *Journal of Economic Behavior & Organization* 108, 59–77.
- Renner, E. and J.-R. Tyran (2004). Price rigidity in customer markets. *Journal of Economic Behavior & Organization* 55(4), 575–593.
- Syverson, C. (2011). What determines productivity? *Journal of Economic Literature* 49(2), 326–365.
- Yang, H. (2013). Nonstationary relational contracts with adverse selection. *International Economic Review* 54(2), 525–547.
- Zaheer, A. and N. Venkatraman (1995). Relational governance as an interorganizational strategy: An empirical test of the role of trust in economic exchange. *Strategic management journal* 16(5), 373–392.

Appendix

A Infinitely Repeated Games

We now provide theoretical support for the main hypotheses of the paper. In this section, we use a standard repeated games approach. In the next section, we will consider a finite horizon game where we allow for committed fairness types.

Here we consider an infinite horizon version of the game.²⁵ The common discount factor is $\delta \in (0, 1)$. $S(x) \equiv 10x - 10$ is the surplus increase that the buyer receives from a seller who chooses x rather than the minimal quality 1. To simplify, we allow for continuum choice sets $[0, 100]$ and $[1, 10]$ for p and x , respectively.²⁶

A.1 Complete Information

We first consider the complete information game. We focus on equilibria in *trigger strategy profiles* (TSP) such that the buyer (seller) chooses some fixed p^* (x^*) as long as neither player has deviated from his assigned choice; after any deviation, both players choose their respective minimal action. A *trigger strategy equilibrium* (TSE) is a subgame-perfect equilibrium in trigger strategy profiles.²⁷

1. Consider the repeated trust contract game with complete information for $\theta \in \{L, H\}$. (p^*, x^*) is sustainable as a TSE for (δ, θ) if and only if

$$p^* \geq c(x^*, \theta) / \delta, \quad (1)$$

$$S(x^*) \geq p^*. \quad (2)$$

The proof is in appendix C. Condition (1) guarantees that the seller who expects a discounted future equilibrium payment δp^* is willing to incur the equilibrium costs of $c(x^*, \theta)$ rather than deviate to the minimal quality. Condition (2) guarantees that the buyer is willing to pay the price p^* rather than zero: She must expect an equilibrium quality x^* generating benefits $S(x^*)$ that are high enough to compensate for the prices p^* .

For $\theta \in \{H, L\}$, all quality levels in $[1, 10]$ are sustainable for suitable prices and sufficiently high discount factors: This follows from (1) and (2) because $S(x) >$

²⁵We ignore the first, cheap-talk, stage at the beginning of the game, as the analysis of the repeated game does not depend on the signal.

²⁶We define the cost function at non-integer values by linear interpolation.

²⁷The results are very similar if we use forgiving strategies where both players only punish once or cut-off strategies where players only punish below a certain level of the other player's action. Details are available upon request.




Figure A.5: Trigger Strategy Equilibrium

$c(x, H) \geq c(x, L)$ for all $x \in [1, 10]$. Figure A.5 illustrates the set of sustainable price-quality vectors for a given parameter choice in the trust contract game.²⁸ The following corollary is immediately intuitive from this figure.

1. *In the repeated trust contract game, fix $\delta > c(10, H) / S(10)$.²⁹*

(i) *For any $\theta \in \{L, H\}$, the lower and upper bound of the set of sustainable x^* are both increasing in p^* .*

(ii) *For any $(\theta, p^*) \in \{L, H\} \times [0, 100]$, the maximal sustainable x^* is weakly higher for $\theta = L$ than for $\theta = H$.*

(iii) *For any $(\theta, x^*) \in \{L, H\} \times (1, 10]$, the minimal p^* for which x^* is sustainable is higher for $\theta = H$ than for $\theta = L$.*

Part (i) of Corollary 1 provides support for Hypothesis 2 which states that we should expect higher quality and higher price going hand in hand. Parts (ii) and (iii) of Corollary 1 both state in slightly different ways that, consistent with Hypothesis 3(i), high quality is harder to sustain with high costs than with low costs: For any given price, the maximal equilibrium quality is higher in the latter case than in the former; similarly, for any given quality, the minimal required price is lower. Finally, the second part of Hypothesis 3 reflects the fact that, according to Condition (1), the lower bound of the equilibrium set is proportional to the costs of producing the respective output. Thus, for high marginal costs this bound is steeper than for low marginal costs. Intuitively, for any given quality increase, the seller requires a greater compensation with high costs

²⁸We fixed $\delta = 0.8$.

²⁹This condition guarantees that (1) and (2) both hold for some prices. In our parametrization, it corresponds to $\delta \geq 0.42$.

than with low costs.

A.2 Incomplete Information

We now proceed to the incomplete information case. We confine ourselves to the sub-game starting after the signal choice of the seller. We suppose that the seller has type L with probability $\mu \in (0, 1)$; and this probability is common knowledge. We think of this probability as the exogenous probability corresponding to the type distribution. We consider pooling strategy profiles in the incomplete information game that correspond to those used in the complete information case. A *pooling trigger strategy equilibrium* (PTSE) is a weakly perfect Bayesian equilibrium in which the players use trigger strategies.

2. *In the trust contract game with incomplete information, (p^*, x^*) is sustainable as a PTSE given δ if and only if (2) and (1) with $\theta = H$ hold.*

Intuitively, in a pooling equilibrium, the binding incentive constraint is that the high-cost types are willing to provide the desired quality. The observed average quality is the same with incomplete information and with complete information. This provides a first way to support Hypothesis 1.

B Fairness

In the following, we discuss Hypothesis 1 in a setting with finite horizon when some players are commitment types à la Kreps et al. (1982). Specifically, the analysis relies on the idea that some players have social preferences. We assume that, for both buyers and sellers, a fraction $\phi \in (0, 1)$ are fair types (F), whereas the remaining $1 - \phi$ are selfish types (S) who maximize monetary payoffs. A fair type responds to the previous action of the other player by choosing what he considers a fair response, as long as he does not put probability 1 onto the other player being a selfish type, in which case he takes the minimal action. There is (two-sided) asymmetric information about fairness types. Thus even the game with complete cost information becomes a game of two-sided asymmetric information, and the game with one-sided incomplete cost information becomes a game with two-sided asymmetric information with two-dimensional seller types: Buyers have types S (selfish) and F (fair). Sellers have types SH, SL, FH and FL, where the first letter corresponds to the preference type (selfish or fair) and the second letter corresponds to the cost type (high or low).

B.1 Complete Cost Information

First, consider the game with complete cost information with cost type $\theta \in \{H, L\}$. In the following, let $p_\theta(x)$ be the price that splits the surplus $10x - 10 - C(x, \theta)$ evenly for the quality level x , so that $10x - 10 - p_\theta(x) = p_\theta(x) - C(x, \theta)$. Similarly, we write $x_\theta(p)$ for the inverse function, that is, the quality that splits the surplus equally given the price p .³⁰ We assume that fair sellers behave as follows:

Assumption A1: *When fair types are sure they are facing a selfish type, they choose the minimal action. Otherwise they behave as follows:*

- (i) *Fair sellers with cost type θ respond to a price p by choosing $x_\theta(p)$.*
- (ii) *Fair buyers start by choosing $p_\theta(10)$. Thereafter, they respond to any quality level x by choosing $p_\theta(x)$ in the following period.*

Fair types are thus committed to non-strategic surplus-splitting behavior, unless they are sure they are facing a selfish type. Intuitively, selfish types can benefit from imitating the fair types and not revealing their selfishness. To confirm this intuition, we investigate the conditions under which an *efficient pooling equilibrium* with the following strategies and beliefs exists:

1. Fair types choose actions as described in Assumption A1.
2. In all periods except the last one, selfish types choose the same actions as fair types, unless they have previously deviated. In the last period, they choose their minimal action.
3. Beliefs about fairness types are given by priors unless the other player has chosen at least one action that a fair player would not choose. In the latter case, beliefs are that the other player is a selfish type with probability one.

If such an equilibrium exists, it involves efficient quality levels, except in the last period for selfish players. Moreover, the price is always $p_\theta(10)$.

3. *Suppose A1 holds. An efficient pooling equilibrium exists if and only if $\phi > \frac{p_\theta(10)}{90}$. As $p_L(10) = 53$ and $p_H(10) = 66$, the required shares of fair players are 0.59 and 0.73 for $\theta = L$ and H , respectively.*

The essential part of the proof (see appendix C) is to make sure that selfish buyers want to pool, even in the final period. This requires a sufficient share of fair players.

³⁰If the price is so high that the maximal quality $x = 10$ leaves at least half the surplus to the seller, then $x_\theta(p) = 10$.

B.2 Asymmetric Cost Information

When there is incomplete information about cost types, it is not immediately clear what "fair prices" are, even when buyers and sellers intend to split the surplus fifty-fifty. In particular, players' fairness assessments could depend on their informational situation, resulting in disagreements. Moreover, one might argue whether fair behavior implies truth-telling or not. We consider two different settings, with different assumptions in these two dimensions.

B.2.1 Setting 1: Pure quality commitment and agreement about fair behavior

We first consider a setting where (1) fair types are not committed to truth telling and (2) both players agree on what constitutes a fair price response $p^f(x)$ to a given quality level x ; with the inverse function $x^f(p)$ capturing the quality that is a fair response to a price p .

With respect to (2), we have two polar cases. First, we consider *credulous* buyers who take a high signal for granted and consider $p^f(x) = p_H(x)$, the fairness price corresponding to high-cost agents under complete information, as an adequate price. Second, we consider *skeptical* buyers who do not think a signal is informative and consider $p^f(x) = \lambda p_L(x_{t-1}) + (1 - \lambda) p_H(x_{t-1})$, the expected fairness price corresponding to the share of each cost type in the population, as an adequate price. In each case, we assume that sellers have the same understanding of fairness, as will be detailed in the following assumptions that apply to credulous and skeptical buyers.

Assumption A2: (i) *When fair buyers are sure they are facing a selfish type, they choose the minimal price.*

(ii) *When fair buyers assign the prior probability ϕ to the seller being fair and probability 1 to cost type θ , they choose $p_\theta(x)$ as in the complete cost information case θ .*

(iii) *When fair buyers assign probability ϕ to the seller being fair and λ to the cost type being low, they choose price $p_t(x_{t-1}) = p^f(x_{t-1})$; in period 1, they choose price $p^f(10)$.*

The assumptions on fair sellers are similar.

Assumption A3: (i) *When fair sellers are sure they are facing a selfish type, they choose the minimal action. They are indifferent about the type they state.*

(ii) *When fair sellers assign probability ϕ to the buyer being fair and expect that the buyer assigns probability 1 to cost type θ , they choose $x_t(p_t) = x_\theta(p_t)$ as in the complete cost information case with cost type θ .*

(iii) *When fair sellers assign probability ϕ to the buyer being fair and expect that the buyer assigns probability $\lambda \in (0, 1)$ to the cost type being low and probability $\phi \in (0, 1]$ to the seller being fair, they choose actions $x_t(p_t) = x^f(p_t)$.*

We are interested in the existence of an *efficient pooling equilibrium* as defined in Section 2.1 (except that Assumption A1 is replaced with Assumption A2 and A3). In such an equilibrium, selfish players pool with fair players (who play A2 and A3), except that selfish sellers choose the minimal quality in the final period; beliefs are given by priors. Beliefs correspond to prior as long as only equilibrium play is observed; otherwise players assume the other player is selfish for sure.

4. *Suppose A2 and A3 hold. There exists an efficient pooling equilibrium if and only if $\phi > \frac{p^f(10)}{90}$, no matter whether buyers are credulous or skeptical.*

While Proposition 4 applies to credulous as well as skeptical buyers, the implications are very different in the two cases. For credulous buyers, $p^f(10) = p_H(10) = 66$. Thus, the condition in the proposition is identical to the one in Proposition 3. Intuitively, as the buyer is prepared to accept the high signal as if it were verifiable, everything is as under complete information. In this case, Proposition 4 thus provides no reason to expect that incomplete information has an effect on quality.

With skeptical buyers, $p^f(10) \in (p_L(10), p_H(10)) = (53, 66)$. Thus, the condition for the efficient pooling equilibrium in Proposition 4 is easier to satisfy than the condition in Proposition 3: Sellers accept a lower price, because they accept the buyer's uncertainty about costs. In this sense, one might conjecture that incomplete information increases quality.

B.2.2 Disagreement on fair behavior; truth-telling commitment

We now modify the set-up of the previous section. First, we modify Assumption A3 by assuming that the fair seller is also committed to signalling the true type (Assumption A3'). Second, we allow for disagreements on what constitutes fair behavior that reflect the informational situation of the players. We will be particularly interested in an equilibrium which separates between type FL and the three remaining types.

Suppose that a buyer is sure he is facing FH, SL or SH, with relative probabilities given by priors. Denote the share of fair types FH in the pool consisting of FH, SL and SH as σ . Moreover, we denote the share of low-cost types in this pool as λ . For our purposes, it suffices to modify Assumption A2 on fair buyers slightly.

Assumption A2': (i), (ii): as in A2.

(iii) *When fair buyers assign probability σ to the seller being fair and λ to the cost type being low, they choose price $p^B(x) = \gamma p_H(x) + (1 - \gamma) p_L(x)$ for some $\gamma \in (0, 1)$; in period 1, they choose price $p_1 = p^B(10) = \gamma p_H(10) + (1 - \gamma) p_L(10)$.*

Part (iii) of the assumption states that the buyer considers a price as fair that is lower than if she was sure that the cost type was low, but higher than if she was sure that the cost type was high. The assumptions on fair sellers are similar.

Assumption A3’: (i),(ii):as in A3.

(iii) *Fair sellers always truthfully reveal their type.*

(iv) *When fair sellers assign probability σ to the buyer being fair and expect that the buyer assigns probability $\lambda \in (0, 1)$ to the cost type being low and $\phi \in (0, 1]$ to the seller being fair, they choose actions $x(p)$ such that the inverse function satisfies $p^S(x) = \delta p_H(x) + (1 - \delta) p_L(x)$ for some $\delta \in (0, 1)$ such that $\delta \geq \gamma$.*

Part (iii) extends the commitment of the fair player to truth-telling. Part (iv) states that, like the buyer, the seller considers a price as fair that is lower than if the buyer was sure that the cost type was low, but higher than if the buyer was sure that the cost type was high. Crucially, however, whenever $\delta > \gamma$, he thinks that the price should be closer to $p_H(x)$ than the buyer does. As a result, her response to a given price is lower than what the seller perceives to be a fair quality.

Thus, our assumptions reflect the idea that under asymmetric information about cost types both parties have different assessments about fairness. The uninformed party (the buyer) thinks that the adequate price lies somewhere between those corresponding to the low cost and high cost cases with complete information. In principle, the informed party acknowledges this (she understands the skepticism of the buyer to some extent), but, when she has high costs, she considers a higher price as adequate than the seller does.

We now show that the asymmetry in the fairness notions can yield a gradual unravelling of cooperation.

For any combination of γ and δ , we can recursively define a sequence $(p, q)_{\gamma, \delta}$ of prices and qualities as follows. Figure B.6 illustrates the construction.³¹

p_1 is given as in Assumption A2’(iii). For all $t \in \{2, \dots, 15\}$, p_t is given as $p_t = p^B(x_{t-1})$, the fair response from the perspective of the uninformed buyer. For all $t \in \{1, \dots, 15\}$, x_t is given by the requirement that $p_t = p^S(x_t)$, the fair response from the perspective of the informed seller. By construction, this sequence has the following properties.

1. In the boundary case that $\gamma = \delta = 1$, then $(p, q)_{\gamma, \delta}$ is constant (except in the last period), given by a price of 66 and a quality of 10.
2. If $\gamma < \delta < 1$, then $(p, q)_{\gamma, \delta}$ is such that both the price and the quality are decreasing over time.
3. $(p, q)_{\gamma, \delta}$ is continuous in γ, δ .

³¹For clarity of presentation, the figure corresponds to the extreme case that $\gamma = 0$ and $\delta = 1$; otherwise the two price lines would be closer together.




Figure B.6: Breakdown of corporation

Note that $\gamma = \delta = 1$ corresponds to the case of a credulous buyer without disagreement on fair prices. Thus, in this case the only difference to the previous section is the truth-telling commitment of the fair seller.

We define a *semi-separating equilibrium with decreasing cooperation* as follows.

- (i) Social players choose their actions as described in Assumptions A2' and A3'.
- (ii) Seller types SL and SH both pool with FH , except in the last period.
- (iii) The selfish buyer pools with the fair buyer.
- (iv) After any history with an initial signal of L , constant prices $p_L(10) = 53$ and qualities 10, the buyer believes that the seller is FL with probability 1.

After any sequence with an initial signal of H such that every buyer price satisfies $p_t = p^B(x_{t-1})$ and every subsequent quality choice satisfies $p^S(x_t) = p_t$, the buyer believes that the seller is one of the three types SH , SL or FH , with relative probabilities given by priors.

In all other cases, the buyer thinks the seller is either type SH or SL for sure.

- (vi) After any history where the seller signaled L and the buyer chooses $p_L(10) = 53$ (and the seller chooses 10) in all periods, the seller maintains prior beliefs about the fairness type of the buyer.

After any history where the seller signaled H and the buyer chooses p_1 defined as in Assumption 2'(iii) and $p_t = p^B(x_{t-1})$ for $t \in \{2, \dots, 14\}$ the seller maintains prior beliefs about the fairness type of the buyer.

After all other histories, the seller assigns probability 1 to the buyer being selfish.

We now provide sufficient conditions for such an equilibrium to exist.

5. Suppose that Assumptions A2' and A3' hold with $\gamma < \delta < 1$, ϕ sufficiently large

and γ sufficiently close to δ . Then the game has a semi-separating equilibrium with decreasing cooperation.

The proof (see below) captures the following intuition. First consider the behavior of fair types after a high signal. Fair buyers start with the price that she deems adequate for the efficient for the efficient effort level. Fair sellers do not consider this price as sufficient (because $\gamma < \delta$); they thus respond by choosing lower effort levels than efficient. In turn, the buyers do not consider these qualities as adequate to justify the original price. They thus respond by reducing the price even more, thereby inducing a further quality reduction. Thus, a further quality reduction takes place. Arguing this way, quality and price decrease over time. To investigate the pooling incentives of selfish sellers, it is most important to compare their payoffs in the equilibrium with the alternative where they imitate the fair L type. In the latter case, they carry out the efficient action, but obtain only the compensation designed for the L type (53). In the former case, that is, in equilibrium, they start out with a compensation that is higher (between 53 and 66) for an effort that is lower than 10. Thus they are initially better off. Over time, the compensation decreases. However, if γ is sufficiently close to δ , however, the decrease is sufficiently slow that deviation is never worthwhile.

The analysis shows that for a wide range of fairness types, including those where the efficient pooling equilibrium exists under complete information, the semi-separating equilibrium with breakdown of cooperation exists. This provides support for the idea that incomplete information may lead to the breakdown of cooperation.

C Proofs

C.1 Proof of Proposition 1

Fix a pooling trigger equilibrium corresponding to (p^*, x^*) . The expected payoffs in periods after histories that trigger p^* are given by $S(x^*) + 10 - p^*$ for the buyer. The optimal deviation after any such history is to give a zero price, yielding payoffs 10. Condition (2) guarantees that this deviation is not profitable. For the seller, the optimal downward deviation is to choose zero quality levels, resulting in cost savings of $c(x^*, H)$ and a payoff reduction of p^* , discounted by δ . Condition (1) with $\theta = H$ guarantees that the deviation is not profitable for seller H . As $c(x^*, H) > c(x^*, L)$, the corresponding condition for seller L is implied. Upward deviations are clearly not profitable either.

After histories triggering the minimal price, the buyer will respond by choosing 0 in the future independent of the seller's current behavior. Thus exerting a non-minimal quality is costly for the buyer, without any corresponding benefits. The argument after histories triggering minimal quality is analogous.

C.2 Proof of Proposition 3

A selfish player who has seen an action of a player that a fair player would not choose is sure that this player is selfish and thus best responds by minimal actions. Thus, we only consider behavior after actions that are consistent with fair types.

Given the proposed beliefs and actions of the buyers, selfish sellers who imitate the fair types earn payoffs $p_\theta(10) - C(10, \theta)$ in all periods except the last one and $p_\theta(10)$ in the last one. Deviation would lead to payoff 0 in every period. As $p_\theta(10) - C(10, \theta) > 0$, imitation is profitable in all periods.

Given the proposed beliefs and actions of the sellers, selfish buyers who imitate the fair types earn payoffs $100 - p_\theta(10)$ in all periods except the last period. In the last period, they obtain a payoff of $100 - p_\theta(10)$ if the seller is fair and $10 - p_\theta(10)$ if the seller is not. Thus the expected payoff is $90\phi + 10 - p_\theta(10)$ in the last period if there has been no previous deviation. Any deviation of a buyer would reveal her type; thus the best deviation is to a price of zero. This would lead to payoff 10 in every period. Thus, imitation is profitable even in the last period if and only if $\phi > \frac{p_\theta(10)}{90}$.

According to the suggested strategies, all players choose the fair actions (except the selfish sellers in the last period). Thus, Bayes' Law implies that beliefs should correspond to priors after fair actions. Also, any deviation from such actions is off equilibrium, so that any posteriors are consistent with Bayes' Law.³²

³²Also, if one argues that selfish sellers literally *cannot* choose anything except their commitment type,

C.3 Proof of Proposition 4

Note first that signals in the quality commitment are pure cheap talk and do not reveal anything about the type.

Given the beliefs of the buyers, selfish sellers who imitate the fair types earn payoffs $p^f(10) - C(10, \theta)$ in all periods except the last one and $p^f(10)$ in the last period. Deviation would lead to payoff 0 in every one. As $p^f(10) - C(10, \theta) > 0$, imitation is profitable in all periods.

Given the beliefs of the sellers, selfish buyers who imitate the fair types earn payoffs $100 - p^f(10)$ in all periods except the last period. In the last period, they obtain a payoff of $100 - p^f(10)$ if the seller is fair and $10 - p^f(10)$ if the seller is not. Thus the expected payoff is $90\phi + 10 - p^f(10)$ in the last period. Any deviation would reveal the type; thus the best deviation is to a price of zero. This would lead to payoff 10 in every period. Imitation is thus profitable even in the last period if and only if $\phi > \frac{p^f(10)}{90}$. A player who has seen an action of a player that a fair player would not choose is sure that this player is selfish and thus best responds by minimal actions.

According to the suggested strategies, all players choose the fair actions (except the selfish sellers in the last period). Thus, Bayes' Law implies that beliefs should correspond to priors after fair actions. Also, any deviation from such actions is off equilibrium, so that any posteriors are consistent with Bayes' Law.

C.4 Proof of Proposition 5

First, consider buyer beliefs. Only seller type FL signals L and chooses 10 thereafter. Thus, it is consistent with Bayes' Law that the buyer believes that the seller is FL with probability 1 after such histories. Similarly, only seller types FH, SH and SL signal H and choose according to $p^S(x_t) = p_t$ thereafter. Thus, it is consistent with Bayes' Law that the buyer believes that the seller is FH, SH or SL with probability 1, with relative probabilities given by priors. All other histories at which the buyer moves are off equilibrium, so that beliefs are not restricted.

Second, consider seller beliefs. As selfish buyers pool with fair buyers, it is Bayes-consistent for them to maintain priors if they have signaled high and then choose according to $p^S(x_t) = p_t$ thereafter and buyers chose prices p_1 as in Assumption 2(iii) and $p_t = p^B(x_t)$ for $t \geq 2$.

Third, consider the behavior of selfish buyers. They earn a positive payoff by pooling, which they cannot by deviating (and thereby revealing their selfish type).

Fourth, consider the behavior of selfish sellers. In period $t = 15$, they obviously earn the maximum possible payoff by choosing the minimal action. Consider any pooling

then a deviation is indeed fully revealing of the selfish type.

history up to period $t \leq 14$. By continuing to pool thereafter, these sellers attain payoffs of $\sum_{\tau=t+1}^{14} (p_t - C(x_t, \theta)) + p_{15}$ thereafter. By revealing their type through deviation in period $t + 1$, they obtain payoffs of 0 in each future period. As $p_t > C(x_t, \theta)$ by construction, a deviation that reveals the type is not profitable in any period. However, the seller could instead deviate by signaling L and then imitating type FL . This yields a payoff of $53 - C(10, \theta)$ in each period (except the last one, where the payoff would be 53). As $p_1 > 53$ by construction, the net profit in period 1 after this deviation is never higher than with the proposed equilibrium strategy. If γ converges to δ , disagreements in fairness assessments are small. Hence, p_t falls only very slowly and the deviation profit in any of the remaining periods is not higher than under the proposed equilibrium strategy. Therefore deviation is not profitable.³³

³³ A more general sufficient (but still not necessary) condition for this equilibrium to emerge is that $p_{15} \geq 53$. It is possible to calculate under which conditions this holds.