HOW DO INFORMAL AGREEMENTS AND REVISION SHAPE CONTRACTUAL REFERENCE POINTS?

Ernst Fehr
University of Zurich

Oliver Hart
Harvard University

Christian Zehnder
University of Lausanne

Abstract
The notion of contracts as reference points provides the basis for a deeper understanding of important phenomena such as the employment contract, vertical integration, firm scope, authority, and delegation. Previous experiments lend support to this notion but they ignore realistic aspects of trading relationships such as informal agreements and ex-post renegotiation or revision. Here we show that the central behavioral mechanism underlying contractual reference points is robust to such considerations. Our data reveal that informal agreements can mitigate the trade-off between rigidity and flexibility but they do not fully resolve the problem of misaligned reference points. Our experiments also show that contract revision is a more nuanced process than the previous literature has recognized. We find, for example, that it is sometimes better for parties to write a simple (rigid) contract and then revise it ex post if needed, rather than to anticipate and include future contingencies in a (flexible) contract from the outset. (JEL: C90, D03, D86)

1. Introduction

A series of recent papers develops the notion that ex-ante contracts serve as reference points for ex-post trade (see Hart and Moore 2008; Hart 2009; Hart and Holmström 2010). The idea is that an initial contract circumscribes what parties feel entitled to, but may not pin down a unique "fair" outcome. In the simplest version each trading party has a self-serving bias that leads him to hope for the most favorable outcome permitted by the contract. In flexible contracts that allow for multiple outcomes, buyers and sellers typically have different favorite outcomes so that misaligned reference points emerge. Although the parties will end up compromising, if they feel that they have not received

The editor in charge of this paper was Stefano DellaVigna.

We gratefully acknowledge financial support from the US National Science Foundation through the National Bureau for Economic Research and the Swiss National Science Foundation (Grant Nos. 100018_135140 and 100018_140734/1). We would like to thank Heski Bar-Isaac, Roland Benabou, Ryan Bubb, Sylvain Chassang, Bob Gibbons, Holger Herz, Louis Kaplow, Wolfgang Pesendorfer, Arno Riedl, Klaus Schmidt, Kathy Spier, Roberto Weber, Philipp Weinschenk, the referees, and the editor for helpful comments.

E-mail: ernst.fehr@econ.uzh.ch (Fehr); ohart@harvard.edu (Hart); christian.zehnder@unil.ch (Zehnder)

DOI: 10.1111/jeea.12098

what they are entitled to, they will be aggrieved and engage in counterproductive ex-post behavior (e.g., quality shading). A rigid contract avoids this situation by limiting the number of outcomes and thereby aligning reference points. The downside of rigidity is that the terms cannot be adjusted to the realized state of the world, which can lead to ex-post inefficiency. The theory implies that there is a trade-off between contractual rigidity and flexibility.

Contractual reference points can explain employment contracts, which fix wages in advance and leave task discretion to the employer (Hart and Moore 2008); indexation in contracts and the role of payoff uncertainty for vertical integration (Hart 2009); and firm scope, authority, and delegation (Hart and Holmström 2010). Initial evidence provides support for the approach (Fehr, Hart, and Zehnder 2009, 2011; henceforth FHZ), but previous studies neglect two key features of real-life trading relationships: informal agreements and ex-post renegotiation or revision.

Informal agreements and ex-post revision can both potentially undermine the theory. To start with the first, the theory assumes that states of the world, although not verifiable, are observable. Observability suggests that trading parties could reach informal, state-contingent agreements: for example, they could agree to split the ex-post surplus 50:50. If the parties can use informal agreements to "manage" expectations, they may be able to achieve the benefits of flexibility without incurring any costs. We conduct experiments that show that although informal agreements mitigate the trade-off between rigidity and flexibility they do not fully resolve the problem of misaligned reference points.

Consider next ex-post revision. Suppose that a buyer and a seller write a rigid contract to align reference points, and ex post the terms of the contract are such that an inefficient outcome will occur. Obviously, if revision is possible, the parties can change the contract to avoid this inefficiency. But this suggests that a rigid contract may not be costly after all: a rigid contract can always be turned into a flexible contract. There is also a more subtle point. Suppose that the rigid contract leads to an ex-post *efficient* outcome and so revision is *not* required. The fact that revision is possible may nonetheless influence parties' entitlements: a party may feel entitled to an outcome outside the initial contract given that the contract can be revised. At an extreme the initial contract might cease to be a reference point altogether. In this paper we show that contracts continue to be reference points in the presence of revision. We also identify circumstances in which a rigid contract that is revised can achieve the benefits of flexibility without incurring its costs.

Our setup builds on the baseline condition of FHZ (2011) (see Figure 1 for a timeline). A buyer and a seller can trade a widget. The buyer's value is known in advance but there is ex-ante uncertainty about whether the seller's cost is high or low. Ex post, trade is voluntary. Ex ante, the buyer can choose between a rigid contract specifying a single price and a flexible contract specifying a price range. Contract terms are determined in a competitive auction among sellers. In rigid contracts the auction directly determines the fixed price; in flexible contracts the auction determines the lower bound of the price range. After the auction, a random device determines whether the seller's cost is high or low. In flexible contracts the buyer can assure trade

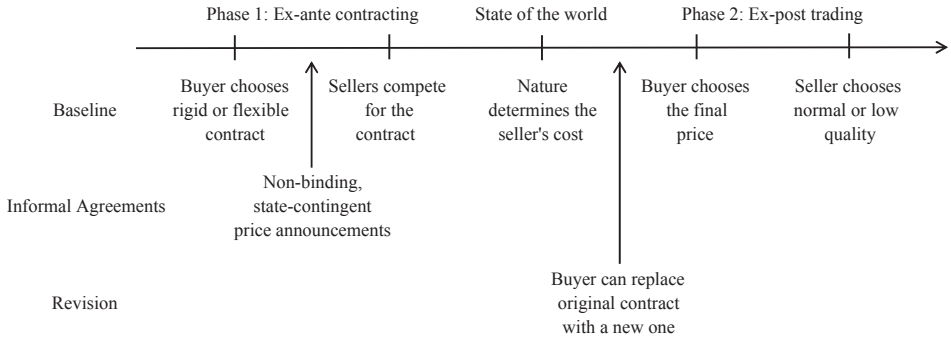

FIGURE 1. Timeline.

by choosing a price from the price range that equals or exceeds the seller's cost. In rigid contracts the price is fixed and cannot be adjusted; this implies that trade is feasible only if the seller's cost is low. If trade occurs, the seller can decide to provide either normal or low quality (i.e., shade). The provision of low quality is slightly costly.

FHZ (2011) confirm Hart and Moore's (2008) prediction that there is a trade-off between contractual rigidity and flexibility. Flexible contracts are useful because they guarantee trade in both states, but flexibility also causes a lot of shading since buyers and sellers seem to have misaligned reference points concerning the final price. Contractual rigidity helps to reduce ex-post shading substantially, because a competitively determined fixed price seems to align ex-ante expectations. But rigid contracts prevent trade from occurring when the cost is high. These results are reassuring for the theory, because most organizational implications of the model follow from the existence of this trade-off.[1]

In this paper we implement two new conditions. In the informal agreement condition buyers have the opportunity to communicate informally their pricing plans in flexible contracts (i.e., they can make nonbinding state-contingent price announcements, which may align the trading parties' expectations). Our data reveal that having this opportunity indeed reduces the shading rate in flexible contracts. As a consequence, flexible contracts become more attractive and are chosen more frequently than in the baseline condition. However, the decrease in shading is moderate and does not eliminate the trade-off between rigidity and flexibility. When informal agreements are available, low prices still trigger more shading in flexible contracts than in rigid ones. As a result, rigid contracts yield higher profits for buyers in the low-cost state than flexible contracts. This advantage is large enough to offset the disadvantage that rigid contracts do not allow for trade in the high-cost state. Even if informal agreements

1. It is also noteworthy that the observed behavior cannot be explained either by traditional contract theory or by standard behavioral models. In particular, existing theories of social preferences (Fehr and Schmidt 1999; Bolton and Ockenfels 2000; Rabin 1993; Charness and Rabin 2002; Dufwenberg and Kirchsteiger 2004; Falk and Fischbacher 2006) cannot account for the trade-off. See FHZ (2011) for a more detailed discussion.

are available, flexible contracts are, on average, no more profitable than rigid contracts. The result that informal agreements do not eliminate the trade-off between rigidity and flexibility is important, especially because the simplicity of our setup (only two states, complete information) gives informal agreements a good chance to be effective.

In our second new condition we allow for ex-post revision. Specifically, we suppose that the buyer has the right unilaterally to replace the existing contract with a new one. Since the seller has no veto, this is actually closer to what lawyers call a "repudiation". We chose this particular form of revision because it provides a powerful stress test for the relevance of contractual reference points: the easier it is to change a contract the less likely it is to serve as a reference point. We find that revision opportunities do not render contractual reference points irrelevant. Although revision is always feasible, the parties do not seem to hope for outcomes outside the ex-ante contract when trade is feasible within the contract. Specifically, if buyers in rigid contracts decide to stick to the agreed-upon price in the low-cost state, the shading rate remains the same as in the baseline treatment. This is a strong finding. As contracts can be changed at no cost, the contract choice is ultimately a framing decision without consequences for feasible outcomes. Nevertheless, sellers seem to accept the competitively negotiated fixed price as a reference point and do not feel entitled to an upward revision of the price. This makes it likely that contractual reference points also remain important when revision is more difficult and/or costly.

Revision improves rigid contracts if costs are high, because revision allows the buyers to increase the price to cover the seller's cost. While these mutually beneficial revisions trigger some shading (probably because of misaligned entitlements caused by the newly introduced flexibility), the gains from trade are still substantial and comparable to those in a flexible contract. Thus, buyers who choose a rigid contract not only benefit from low prices and low shading rates in the low-cost state; revision allows them to realize the same profits as with flexible contracts in the high-cost state. Therefore, in our revision treatment, rigid contracts yield significantly higher profits than flexible contracts.

However, revision is problematic to the extent that it leads buyers to engage in nonmutually beneficial revisions: specifically, to replace the contract in order to lower the price and grab a larger share of the surplus. We show that such opportunistic revisions tend to cause sellers to engage in substantial shading. This finding is consistent with Hart (2009), who uses a similar view of revision or renegotiation to explain the empirical finding that contracts exhibit a self-enforcing range: they are respected in "normal" states of the world but are breached in exceptional circumstances (see, e.g., Klein 1996).

Perhaps the most important implication of our revision treatment is that our findings provide support for a richer and more realistic view of the revision or renegotiation process. In standard incomplete contracting models the parties can always do better by committing not to renegotiate. As any renegotiation process can be built into the initial contract, any further renegotiation simply adds incentive constraints (see, e.g., Maskin and Moore 1999). Thus we would expect in reality to see parties going out of their way to make ex-post renegotiation difficult. Yet there is little evidence that parties

deliberately put sand in the gears of the renegotiation process.[2] Our study helps to explain this. If the parties build the revision process into the original contract—in our experiment this is equivalent to picking a flexible contract—then this raises the seller's feelings of entitlement in states where revision is not needed (low-cost states) as well as in states where it is (high-cost states). The parties can do better by not incorporating revision initially and then revising ex post as needed.

Our study is related to a number of recent experiments on behavioral contract theory. Brandts et al. (2013) explore the role of contractual reference points in a setup without ex-ante competition and show that informal agreements may have a larger impact in purely bilateral environments. Hoppe and Schmitz (2011), Bartling and Schmidt (2014), and Iyer and Schoar (2012) report evidence consistent with our finding that opportunistic renegotiations are perceived as unfair and trigger lots of shading. We discuss these papers in much more detail after the presentation of our results (see Section 5).

The remainder of the paper is organized as follows. In Section 2, we describe the design of our experiment. Section 3 contains the behavioral predictions. We present our results in Section 4. Section 5 discusses related research and Section 6 concludes.

2. Experimental Design

We present the market setup and the parameters in Section 2.1. Section 2.2 describes the interaction of buyers and sellers. The details of the investigated experimental conditions are provided in Section 2.3. We describe the laboratory procedures in Section 2.4.

2.1. Market Setup and Parameters

Each experimental session has an equal number of buyers and sellers. In every period of the experiment buyers and sellers have the possibility of trading a product. Each seller can sell up to two units, while every buyer can buy at most one unit of the product per period. Hence the supply of the product is twice as large as the demand, and sellers face competition for buyers. When a buyer purchases a unit of the product from a seller, his payoff is equal to his valuation for the product v minus the price p. The payoff of the seller is defined as the difference between the price p and the production cost c. The buyer's valuation for the product depends only on the seller's ex-post quality choice q. The seller's production cost, in contrast, also depends on the realized state of the world σ. There are two states of nature: a good state ($\sigma = g$), in which the seller's production costs are low, and a bad state ($\sigma = b$), in which the production costs are high. The good state occurs with probability $w^g = 0.8$.

The payoffs of buyers and sellers can be summarized as follows:

$$\text{Buyer}'\text{s payoff:} \quad \pi_B = V(q) - p.$$

2. For a contrary view, see Rajan (2012).

TABLE 1. Experimental parameters.

State of nature	Good [Prob($s = g$) = 0.8]		Bad [Prob($s = b$) = 0.2]	
Seller's quality	normal ($q = q^n$)	low ($q = q^l$)	normal ($q = q^n$)	low ($q = q^l$)
Seller's costs	20	25	80	85
Buyer's valuations	140	100	140	100

Notes: The table summarizes the main parameters of the experiment. Buyers' valuations for the product and sellers' production costs are displayed for both states of nature and both quality levels available to the seller.

$$\text{Seller}'\text{s payoff:} \quad \pi_S = p - C(q,\sigma).$$

When trade occurs sellers can choose between two quality levels: normal quality ($q = q^n$) or low quality ($q = q^l$). The production costs for low quality are slightly higher than the production costs for normal quality: $c(q^l, \sigma) > c(q^n, \sigma)$. This reflects the idea that sellers can minimize costs if they simply provide the product desired by the buyer. However, they can sabotage output (at a small cost) if they want to.[3] For each unit of the product which a seller cannot sell—either because he did not manage to conclude a contract with a buyer or because his contract does not allow for a mutually beneficial trade—he realizes an outside option $x_S = 10$. When a buyer is unable to trade, he also realizes an outside option $x_B = 10$. Table 1 summarizes the cost and value parameters of the experiment.

In the experiment sellers and buyers interact in groups of four (two buyers and two sellers). To minimize the role of reputational considerations, these groups are randomly reconstituted at the beginning of each period. Thus, our protocol induces a series of one-shot interactions.

2.2. Interaction of Buyers and Sellers Within a Period of the Experiment

We describe the different steps characterizing the interaction of buyers and sellers in *all* our treatments. Particularities of the different conditions are discussed in the next section.

2.2.1. Random formation of interaction groups. In every period groups consisting of two buyers and two sellers are randomly formed.

3. In Hart and Moore (2008), parties are assumed to be indifferent between shading (sabotaging) and not. We introduce costly sabotage to rule out equilibrium sabotage under standard economic assumptions. (We suppose that sabotage increases rather than decreases costs for reasons explained in FHZ 2011.) The quality choice of the seller in our experiment is similar to costly punishment technologies that have been used in many other cooperation experiments (see, e.g., Fehr and Gaechter 2000 for a typical example). However, our experiment differs from typical gift exchange experiments (see, e.g., Fehr, Goette, and Zehnder 2009 for a review of this literature). In gift-exchange games the pecuniary incentive for workers (i.e., sellers) is to provide the minimal effort (i.e., quality) level, whereas in our paper the normal quality level maximizes seller earnings.

Phase 1: Ex-ante contracting.

Step 1: Buyers' contract choice. Each transaction begins with the buyer's choice of a contract type (t). The buyer chooses between a rigid contract ($t = r$) or a flexible contract ($t = f$). Rigid contracts define a single transaction price p^r ex ante. Flexible contracts, in contrast, specify a price range $[p^l, p^u]$ from which the buyer will choose the price ex post. The buyer can choose only the type of contract, but not the terms. The terms (i.e., the fixed price or the price range, respectively) are determined in a competitive auction among the sellers. In one of our conditions buyers have the option to make informal price announcements if they choose a flexible contract at this stage. In the other conditions no communication possibility is available.

Step 2: Sellers' contract auction. After both buyers in an interaction group have chosen their type of contract, the two contracts are auctioned off to the sellers. The sequence of the auctions is randomly determined within each group. If a rigid contract is auctioned off the auction directly determines the fixed price $p^r \in [c(q^l,g) + x_S, 75] = [35, 75]$.[4] In a flexible contract the auction determines the lower bound of the price range $p^l \in [35, 75]$. The upper bound of the price range is exogenously fixed and equal to the buyer's valuation of the product when the seller provides normal quality: $p^u = \mathrm{v}(q^n) = 140$. Thus, in both cases the auction starts off at 35 and then increases by one unit every half second. Each of the two sellers has a button that allows him to accept the contract at any time during the auction. The first seller who is willing to accept the displayed fixed price or the displayed lower bound, respectively, gets the contract. The seller who loses the auction directly realizes the outside option x_S (i.e., there is an outside option for each of the two auctions).

2.2.2. Determination of the state of the world. After the contract auctions the computer randomly determines the state of the world for each contract independently.

The state is common knowledge to the trading parties.

Phase 2: Ex-post trading.

Step 3: Buyers' choice of contract terms. Once the state has been revealed, the buyer determines the final terms of the contract. How much flexibility he has in doing this depends on the experimental condition and the ex-ante chosen contract. To initiate a mutually beneficial trade the buyer needs to be able to pick a price that covers the seller's cost. (It should be emphasized that trade occurs whenever price covers cost: a seller cannot refuse to trade if price covers cost although he can of course shade, i.e., choose low quality.) The flexible contract always allows for such a choice, but

4. The minimum of 35 for the fixed price ensures that the seller cannot make losses relative to his outside option in the good state even if he provides low quality. This feature guarantees that sellers do not refrain from choosing low quality, just because they want to avoid losses (loss aversion). The maximum of 75 for the fixed price ensures that the price is always below the seller's cost in the bad state of the world. This guarantees that trade is infeasible within rigid contracts if the bad state is realized. However, in the experiment the competitive forces in the auction were strong enough so that the maximum was never binding.

the fixed-price contract does so only in the good state; in the bad state the fixed price of a rigid contract is lower than the seller's cost ($p^r \in [c(q^l,g) + x_S, 75] < c(q^n,b) = 80 < c(q^l,b) = 85$). In the latter case trade is feasible only if the buyer can revise the contract (revision is permitted in only one of our experimental conditions). If the buyer cannot or does not want to revise the contract, trade does not occur and both trading parties realize their outside options. If the contract allows for trade the buyer either pays the fixed price (rigid contract) or picks a price out of the available price range (flexible contract, or revised contract).[5]

Step 4: Sellers' quality choice. Sellers observe the price choice of their buyer and then determine their quality. The sellers always have the choice between normal (q^n) and low (q^l) quality. Remember that choosing low instead of normal quality increases the seller's cost by five units irrespective of the contract type and realized state of the world (see Table 1).

2.2.3. Payoffs and Market Information. When all decisions have been made, profits are calculated and displayed on subjects' screens. In addition, to their profit information buyers also get some aggregated information about the market outcome.[6]

Subsequently, a new period begins and participants are randomly reassigned to new groups.

2.3. Experimental Treatments

In the following, we describe the baseline condition introduced in FHZ (2011) and our two new treatments.

2.3.1. The Baseline Condition (BL). In the baseline condition neither informal agreements in the form of price announcements nor revision opportunities are available.

2.3.2. The Informal Agreement Condition (IA). In the informal agreement condition buyers who choose a flexible contract in Step 1 can decide whether they want to combine the contract with a message of the following form:

> "If costs are low, I plan to pay a price of $p^A(g)$. If costs are high, I plan to pay a price of $p^A(b)$."

5. In the bad state the buyer has to ensure that the price is such that the seller cannot make losses; that is, he must choose a price $p \in [c(q^l,b) + x_S, v(q^n)] = [95, 140]$. Again we do not allow prices to be such that the seller can make losses by choosing low quality, since we want to avoid the possibility that people refrain from shading because of loss aversion (see also footnote 4).

6. The buyers are informed about average payoffs in rigid and flexible contracts of all buyers in all previous periods. In addition, they also learn how many buyers have chosen rigid and flexible contracts in the current period. The aim of the provision of this information is to make learning easier for buyers. Since our setup allows for many possible constellations (two contract types, two states of nature, two quality levels, many prices), learning from individual experience is rather difficult.

The price announcements are in no way binding for the buyer (i.e., the message does not affect the range of actual prices available to the buyer ex post).[7] The buyer can always pick prices which are higher or lower than the announced price if the competitively determined price range of the contract allows this. All market participants are informed about the presence of the message opportunity in flexible contracts in the instructions of the experiment (i.e., the availability of messages is common knowledge in the experiment).

In the informal agreement condition revision of contracts is not permitted. Accordingly, rigid contracts allow for trade only in the good state of the world. In the bad state trading parties with a rigid contract have to realize their outside option.

2.3.3. The Revision Condition (RE). In the revision condition, there are no informal agreements. However, buyers always have the possibility of revising the contract ex post (see previous Step 3). If a buyer decides to revise the contract, the original contract is no longer of relevance and the buyer can choose any price that satisfies $p \in [c(q_l, \sigma) + x_S, 140]$. The seller cannot veto the buyer's decision to revise the terms. Revision is available for rigid and flexible contracts in both states of the world. That is, the buyer can always decide whether he wants to stick to the competitively concluded ex-ante contract and accept the imposed restrictions (i.e., the fixed price in rigid contracts and the lower bound of the price range in flexible contracts, respectively) or whether he wants to abandon the contract and pick his price without restrictions.

It is useful to distinguish three types of revision which may occur in this condition. First, the buyer may revise a rigid contract in the bad state of the world. This allows for a price increase and makes trade feasible. As both parties benefit (at least weakly) from such a revision, we call this a "mutually beneficial revision". Second, the buyer may revise a contract in the good state of the world in order to decrease the price to a level below the ex-ante agreed-upon fixed price or lower bound of the price range, respectively. We call this an "opportunistic revision", because the buyer intends to increase his own profit at the expense of the seller. Finally, there is also the possibility that a buyer voluntarily increases the fixed price of a rigid contract in the good state of the world. We call these costly attempts to increase the seller's profit "altruistic revisions".

2.4. Subjects, Payments, and Procedures

All subjects were students of the University of Zurich or the Swiss Federal Institute of Technology Zurich (ETH). We used the recruitment system ORSEE (Greiner 2004). Each subject participated in only one session. Subjects were randomly subdivided into two groups (buyers and sellers) before the start of the experiment. The subjects' roles remained fixed for the whole session. All interactions of participants were completely

7. This also implies that the buyer cannot stick to his price announcement if the announcement turns out to be below the lower bound of the price range determined in the auction. We explain this in detail to the participants in the instructions.

anonymous. The experiment was programmed and conducted with z-Tree (Fischbacher 2007).

To make sure that subjects fully understood the payoff consequences of the available actions, each subject had to read a detailed set of instructions before the session started. Participants then had to answer several questions about the feasible actions and the payoff consequences of different actions. We started a session only after all subjects had correctly answered all questions. The exchange rate was 15 points = 1 Swiss franc.

In order to make the sellers familiar with the auction procedure we implemented trial auctions before we started the actual experiment. In the trial phase each seller had his own auction (i.e., they did not compete with another seller and no money could be earned).

The data used in this paper were collected in several waves. For the baseline condition we use both the original data of FHZ (2011) (collected in 2007 and 2008) as well as new data collected in 2013. For the informal agreement condition and the revision treatment we use data collected in 2011 and new data collected in 2013. The recruitment procedures, the experimental instructions, the software, and the computer laboratory were identical for all sessions. In addition, participants were always recruited from the same population (but obviously from different cohorts). In the new sessions conducted in 2013 the three treatments were randomly allocated to sessions. These sessions are therefore completely free from any potential confound stemming from intertemporal changes in the subject pool. Since we were unable to detect systematic differences in the results of the different treatments as a function of the data collection period, we pool the data from the different elicitation waves for the analysis. However, we also provide an Online Appendix in which we repeat the analysis only using the data from the sessions conducted in 2013. In the sessions in 2007, 2008, and 2011 we implemented only one matching group per session. In the new sessions in 2013 we increased the number of independent observations by having two matching groups per session. Overall, we have data from twelve independent matching groups for the baseline condition (five sessions from 2007 and 2008 and four new sessions[8]) and eleven independent matching groups for each of the other two conditions (five sessions from 2011 and three sessions from 2013). As a result of no-shows the number of participants per session varies somewhat. In the 16 sessions with one matching group we had 28 participants in nine sessions, 24 participants in six sessions, and 16 participants in one session. In the nine sessions with two matching groups we had 32 participants in seven sessions, and 28 participants in the other two sessions. Thus, in total our analysis is based on data from 692 subjects. A session lasted about two hours and subjects earned on average about 50 Swiss francs (including a show-up fee of 10 Swiss francs).

8. For one of the new sessions in the baseline conditions there was an exceptionally high rate of no-shows. In this session we implemented only one matching group (with 24 participants). To compensate for the loss of data, we conducted another session with two matching groups.

3. Behavioral Predictions

In this section we derive a set of hypotheses for our experiment. In Section 3.1 we present the predictions that result from the assumption that people are purely self-interested money-maximizers. While we do not believe that the self-interest hypothesis is an accurate description of our participants' behavior, we still feel that these predictions are a useful benchmark, not least because much of the theoretical literature on incomplete contracts is based on models that assume pure self-interest. In Section 3.2 we discuss how the presence of contractual reference points affects the predictions for our experiment.

3.1. Predictions under Pure Self-Interest

The prediction of the self-interest model is straightforward. Buyers anticipate that selfish sellers are never willing to engage in costly shading and therefore offer the lowest price permitted by the contract. Competition in the contract auctions implies that the fixed price in rigid contracts and the lower bound in flexible contracts are at the competitive level (i.e., $p^r = 35$ and $p^l = 35$).[9] This implies that rigid and flexible contracts yield the same profit for buyers in the good state of the world ($\pi_B = v(q^n) - p = 140 - 35 = 105$). In the bad state payoffs depend on whether revision is available. If the buyer can revise the contract, both contracts yield the same profit for the buyer ($\pi_B = v(q^n) - p = 140 - 95 = 45$) and the buyer is indifferent between the two. If revision is not possible, the rigid contract results in the outside option ($\pi_B = x_B = 10$) and therefore the buyer strictly prefers the flexible contract. Whether or not informal agreements are available does not affect the predictions.

3.2. Predictions if Contracts are Reference Points

In this section we discuss how the Hart–Moore notion that ex-ante contracts provide reference points for ex-post trade affects the predictions for each of our experimental conditions. For the baseline treatment of FHZ (2011) contractual reference points have the following implications: while contractual reference points do not affect the prediction that the contract auctions yield competitive outcomes ($p^r = p^l = 35$), they change the consequences of the buyers' contract choice. Of particular importance is the fact that flexible contracts may imply that sellers feel entitled to high prices.[10] If a buyer picks a price which is below the seller's reference price, the seller may be

9. Since $p = 35$ corresponds to $p = c(q^l,g) + x_S$ and the seller must offer at least $p = c(q^l,b) + x_S = 95$ in the bad state of the world, a seller can never be worse off if he accepts a contract than if he accepts his outside option.

10. Hart and Moore (2008) assume that each party compares the ex-post outcome to the most favorable outcome permitted by the contract. In FHZ (2011) we extend the model to allow for the case where parties may have heterogeneous reference points (i.e., we take into account that some traders may feel entitled to an outcome other than the most favorable outcome). We show that the predictions of such an extended model remain very similar to those of the original model.

aggrieved and engage in shading.[11] As sellers may have heterogeneous reference points the frequency of shading is predicted to be decreasing in price. Thus, depending on the distribution of sellers' reference prices it can be optimal for the buyer either to avoid shading by increasing the price to a level above the lower bound of the price range or to accept the risk of getting low quality. Rigid contracts should avoid the shading problem, because they pin down the price from the outset and thereby fix expectations. If the shading problem in flexible contracts is severe enough, rigid contracts may be more profitable for buyers, even though they prevent the parties from trading in the bad state of the world.

In the informal agreement condition the central question is whether and to what extent the possibility of making ex-ante price announcements allows buyers to "manage" sellers' reference points. If price announcements give buyers some control over sellers' reference points, buyers should try to push down reference points by announcing the lowest prices possible. Specifically, we would expect buyers to announce the competitive price for the good state ($p^A(g) = 35$) and the lowest feasible price for the bad state ($p^A(b) = 95$). As a consequence, sellers may feel less entitled to high prices and therefore low prices should trigger less shading in flexible contracts as compared to the baseline condition.

In the extreme, informal agreements give buyers full control over seller's reference points (i.e., the seller never feels entitled to prices beyond those announced by the buyer). In this case, buyers would always choose flexible contracts with low price announcements, because such contracts would not only allow them to implement the first best (i.e., shading-free trade in both states), but they would also be able to reap (almost) all the available gains from trade.[12] More realistically, however, informal agreements will give buyers only limited control over sellers' reference prices. The buyer's optimal contract choice will then depend on the base level of shading in flexible contracts and the extent to which informal agreements reduce this shading rate. The larger the impact of low price announcements on sellers' reference points, the more attractive are flexible contracts for buyers.

The revision treatment addresses two important questions regarding contractual reference points. The first one is, do contracts remain reference points even when revision is feasible? And the second question is, how is the reference point determined once revision has been initiated? The first question is of great interest for the cases where the parties decide to stick to the ex-ante agreed-upon contract. Basically, there are two interesting possibilities. One is that the parties continue to hope for their preferred outcome within the limits of the contract (i.e., the contract remains the reference point

11. The assumption that sellers who feel aggrieved are willing to engage in costly shading activities builds on the extensive literature on negative reciprocity. Starting with the seminal paper by Güth, Schmittberger, and Schwarze (1982) on the ultimatum game, behavioral economists have accumulated a great deal of evidence that people who feel treated unfairly are willing to give up material payoffs to punish those who they blame for the unfair outcome (see, e.g., Camerer 2003 for a review of this line of research).

12. As our design ensures that sellers cannot make losses even if they engage in costly shading, a small part of the surplus (five points) would still go to sellers in this equilibrium (see footnotes 4 and 5 for more details).

even though revision would have been possible). The argument for this possibility would be that sellers explicitly agree to the contract in the competitive bargaining process which turns the contract into a focal point that defines their expectations. The alternative view would be that contracts completely lose their meaning in the presence of revision opportunities and the parties simply hope for their preferred outcome within the set of all feasible outcomes (including ones that can be reached only if the contract is revised). This would imply that the reference prices are independent of the contract type, so that the buyer's contract choice would no longer make any difference for ex-post performance. Which of these two views is realistic can be determined only from the data.

Obviously, the second question (i.e., the definition of reference points after the initiation of revision) is important only if contracts continue to shape reference points when revision is not initiated (otherwise the parties just hope for their globally preferred outcome, which is independent of the contract choice). Suppose that contracts *do* continue to shape reference points when revision is not initiated. One possible view is that revision simply turns the existing contract (be it a flexible or a rigid contract) into a completely flexible contract with price range $p \in [c(q_l,g) + x_S,\ v(q^n)]$. We think that this view is a plausible one when it comes to mutually beneficial revision. In our setup mutually beneficial revision occurs when the buyer revises a rigid contract in order to be able to increase the price in the bad state of the world. In this case the situation is indeed very similar to that in a flexible contract: both revised rigid contracts and flexible contracts allow for the same price choices and there is no obvious reason why the seller should respond differently in the two situations. However, we do not think that this view is accurate when it comes to opportunistic revision. In our experiment opportunistic revision can occur if a buyer revises a rigid contract in the good state of the world and picks a price below the competitively determined fixed price, or if a buyer revises a flexible contract in the good state of the world and picks a price below the competitively determined lower bound of the price range.[13] Since the buyer intentionally lowers the seller's profit, opportunistic revision itself may be an important source of aggrievement on the seller side.

4. Results

In this section we present our results. We compare both the informal agreement treatment (IA) and the revision treatment (RE) to the baseline treatment (BL) and discuss additional results that the two new treatments reveal.

13. It is important to mention that neither of these situations can occur in competitive equilibrium. If the fixed price or the lower bound of the price range is at the competitive level, the buyer cannot lower the price after revision. However, since we expect that auction outcomes will often deviate from the competitive level (although they usually converge to the competitive level over time, see FHZ 2011 for details on the baseline condition), it is useful to consider these situations anyway.

4.1. Informal Agreements

In the IA our main interest is to investigate whether the possibility of making nonbinding ex-ante price announcements allows buyers to "manage" sellers' reference points. In Section 3.2 we hypothesize that if announcing prices gives buyers control over sellers' reference points, this should weaken or even eliminate the trade-off between contractual flexibility and rigidity. In particular, we should observe that (i) flexible contracts in the IA are combined with low price announcements, (ii) sellers in flexible contracts engage less often in shading in response to low prices in the IA than in the BL, and (iii) the profitability of flexible contracts relative to rigid contracts is much higher in the IA than in the BL and buyers are therefore much more likely to choose flexible contracts in the IA.

Table 2 and Figure 2 reveal that these predictions are partially supported by the data. Table 2 summarizes the most important outcomes for all three treatments. The table reports average prices, shading rates, average auction outcomes, profits and contract frequencies for rigid and flexible contracts separately. In addition, the table also indicates how often buyers attach a message to a flexible contract in the IA. Figure 2 depicts prices and shading rates in the BL and IA graphically. While average prices in flexible contracts are essentially at the same level in the two treatments (good state: 47.1 (IA), 48.7 (BL) / bad state: 98.1 (IA), 97.6 (BL)), the shading rates are lower in the IA (good state: 0.12, bad state: 0.16) than in the BL (good state: 0.23, bad state: 0.23). However, only the effect for shading in the good state is significant.[14,15] In line with the previous hypotheses this increases the average buyer profit in flexible contracts in the IA (77.8) relative to the BL (72.1).[16] Prices, shading rates, and buyer profits in rigid contracts remain largely unaffected by the possibility of sending informal price announcements in flexible contracts (prices: 38.8 (IA), 40.7 (BL) / shading rates: 0.04 (IA), 0.06 (BL) / buyer profits: 81.8 (IA), 78.8 (BL)).[17] As a consequence, buyers choose the flexible contract more often in the IA (70%) than in the BL (52%).[18]

However, it is important to emphasize that the decrease in the shading rate in flexible contracts in the IA relative to the BL is not pronounced enough to eliminate the trade-off between contractual rigidity and flexibility in the IA. It is still the case that

14. Ranksum tests (flexible contracts, IA versus BL): price: $p = 0.218$ (good state), $p = 0.423$ (bad state); shading: $p = 0.001$ (good state), $p = 0.422$ (bad state).

15. Because we randomly rematch participants within a matching group at the beginning of every period of the experiment, observations stemming from the same matching group cannot be treated as independent. We therefore perform all our nonparametric tests (ranksum, signrank) using one independent observation per matching group (all nonparametric tests are performed using matching group averages of the corresponding variable). All reported p-values for nonparametric tests are based on two-sided tests. All reported standard errors in regression analyses are adjusted for clustering at the matching group level.

16. Ranksum test (flexible contracts, IA versus BL): buyer overall profit: $p = 0.005$.

17. Ranksum tests (rigid contracts, IA versus BL): price: $p = 0.242$; shading: $p = 0.712$ / profits: $p = 0.424$.

18. Ranksum tests (IA versus BL): Relative frequency of flexible contracts: $p = 0.009$.

TABLE 2. Summary statistics—all treatments.

	Baseline (BL)				Informal agreement (IA)				Revision (RE)			
Contract	Rigid		Flexible		Rigid		Flexible		Rigid		Flexible	
State	Good	Bad	Good	Bad	Good	Bad	Good	Bad	Good	Bad	Good	Bad
Message sent	–	–	–	–	–	–	0.52		–	–	–	–
Av. price announcem.	–	–	–	–	–	–	52.0	102.1	–	–	–	–
Revised contracts	–	–	–	–	–	–	–	–	0.09	0.80	0.06	0.00
Average price	40.7	–	48.7	97.6	38.8	–	47.1	98.1	38.7	97.6	43.2	96.9
Rel. freq. of shading	0.06	–	0.23	0.23	0.04	–	0.12	0.16	0.06	0.15	0.13	0.17
Av. auction outcomce	40.5		39.9		38.9		38.4		38.4		38.5	
Average profit buyer	97.0	10.0	82.1	33.2	99.5	10.0	88.1	35.4	99.0	31.1	91.7	36.3
	78.8		72.1		81.8		77.8		85.4		80.8	
Average profit seller	20.4	10.0	27.5	16.4	18.5	10.0	26.5	17.3	18.4	15.5	22.6	16.1
	18.2		25.3		16.8		24.7		17.8		21.3	
Share of contracts	0.48		0.52		0.30		0.70		0.43		0.57	

Notes: The table summarizes the outcomes for rigid and flexible contracts in all treatments. If meaningful the data are provided for each state of the world separately, otherwise the data are reported for both states together (centered entries in the table). "Message sent" indicates the share of contracts in which buyers attached a price announcement to a flexible contract. "Av. price announcem." is the average price announcement made in contracts that had a message attached. "Revised contracts" indicates the share of revised contracts. "Average price" is the average ex-post price paid by buyers. "Rel. freq. of shading" measures the relative frequency of shading. "Av. auction outcome" stands for the average fixed price (rigid contracts) or average lower bound of the price range (flexible contracts) determined in the competitive auctions. "Average profit buyer (seller)" are the average profits for the respective trading party. "Share of contracts" indicates the relative frequency with which contracts have been chosen.

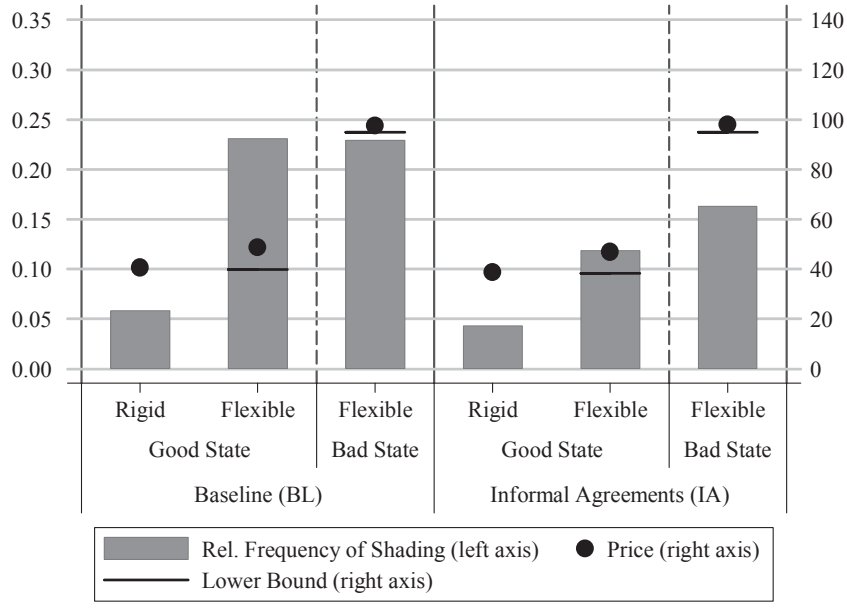

FIGURE 2. Baseline versus informal agreements: prices and shading rates. The figure is based on data from the BL and IA. In the IA flexible contracts include contracts with and without price announcements. Bars represent relative shading frequencies (left axis). Lines display the average lower bound of the price range, and dots represent average actual prices (right axis).

in the good state of the world both prices and shading rates are significantly higher in flexible contracts than in rigid contracts (price: 47.1 (flexible), 38.8 (rigid) / shading: 0.12 (flexible), 0.04 (rigid)).[19] As a consequence, buyers who pick a rigid contract realize significantly higher payoffs (99.5) than those with a flexible contract (88.1) when the good state of the world is realized.[20] The advantage of rigid contracts in the good state of the world is large enough to compensate for the disadvantage that rigid contracts do not allow for trade in the bad state. Overall, buyers with rigid contracts realize an average profit of 81.8, while buyers with flexible contracts make an average profit of 77.8. However, unlike in the BL condition, the difference in overall profits of buyers across contracts is not significant in the IA treatment.[21]

The regression analysis in Table 3 provides further support for our results. Column (1) compares shading rates in rigid and flexible contracts when the good state has been realized. Using observations from the IA and the BL we regress an indicator variable for shading on indicator variables for flexible contracts, the IA, and an interaction of the two. This regression confirms the following results:

19. Signrank tests (IA, rigid versus flexible contracts): price: $p = 0.003$; shading: $p = 0.004$.

20. Signrank tests (IA, rigid versus flexible contracts): buyer profit (good state): $p = 0.003$.

21. Signrank tests (IA and BL, rigid versus flexible contracts): buyer profit (overall): $p = 0.008$ (BL), $p = 0.477$ (IA).

TABLE 3. Shading rates in the BL and the IA.

Dependent variable	Shading BL & IA	Shading BL & IA	Shading BL & IA	Shading BL & IA
Treatment	Good	Good	Bad	Bad
State	(1)	(2)	(3)	(4)
Flexible contract	0.173***	0.259***		
	(0.016)	(0.038)		
IA	−0.015	−0.030	−0.066	−0.068
	(0.015)	(0.020)	(0.059)	(0.071)
Flex. contr. x IA	−0.097***	−0.128***		
	(0.020)	(0.042)		
Price increment		0.000		−0.012***
		(0.001)		(0.004)
Flex. contr. x Price inc.		−0.007***		
		(0.002)		
IA x Price inc.		0.004		0.002
		(0.003)		(0.005)
Flex. contr. x IA x Price inc.		−0.001		
		(0.003)		
Constant	0.058***	0.056***	0.229***	0.259***
	(0.010)	(0.015)	(0.047)	(0.056)
Observations	2826	2826	426	426
Clusters (Matching groups)	23	23	23	23
R^2	0.048	0.072	0.007	0.026

Notes: Columns (1) and (2) use data from rigid and flexible contracts in the BL and the IA in which the good state has been realized. Columns (3) and (4) use data from flexible contracts in the BL and the IA in which the bad state has been realized. "Flexible contract" is an indicator variable which is unity if the contract is of the flexible type and zero otherwise. "IA" is an indicator variable which is unity if the observation comes from the IA and zero otherwise. "Price increment" is equal to price minus 35 (good state) or price minus 95 (bad state). This implies that the constant measures the shading rate in rigid contracts in the BL at the competitive price level (35 or 95, respectively). All columns report coefficients of OLS estimations (linear probability model). Since observations within matching groups may be dependent all reported standard errors are adjusted for clustering at the matching group level.
***Significant at 1%.

(i) The shading rate in rigid contracts is low in the BL (constant = 0.058).

(ii) The shading rate in flexible contracts is significantly higher in the BL ("flexible contract" = 0.173, $p < 0.001$).

(iii) There is no significant difference between shading rates in rigid contracts across treatments ("IA" = −0.015, $p = 0.317$).

(iv) Although the shading rate is still significantly higher in flexible contracts than in rigid contracts in the IA (F-test: "flexible contract" + "flexible contract × IA" = 0.076, $p < 0.001$), the difference is significantly smaller than in the BL ("flexible contract × IA" = −0.097, $p < 0.001$).

In column (2) we additionally control for prices and also interact prices with all other independent variables. Including these additional variables does not change any of the results obtained in column (1). In particular column (2) reinforces again the

finding that the reduction of the shading rate in flexible contracts is not strong enough to eliminate the disadvantage of the flexible contract completely. In other words, in the IA the shading rate at competitive prices remains significantly higher in flexible contracts than in rigid contracts (F-test: "Flexible contract" + "Flex. contr. × IA" = 0.132, $p < 0.001$). Columns (3) and (4) compare shading rates in flexible contracts in which the bad state has been realized across treatments. Column (3) estimates the uncontrolled effect, while column (4) controls for price and also interacts price with the indicator variable for the IA. Both estimations confirm that the shading rates in flexible contracts in which the bad state has been realized are lower in the IA than in the BL, but the differences are not significant (probably because of the smaller number of observations).

We find these results interesting. It seems that informal price announcements allow buyers to reduce shading in flexible contracts, but the effect is not pronounced enough to eliminate the trade-off between rigidity and flexibility.

Table 2 also points out two other interesting results in the IA treatment: (i) about half of the buyers (48%) who pick a flexible contract do not attach an informal price announcement to their contract, and (ii) the average price announcement that buyers attach (if they attach one at all) is much higher (good state: 52.0, bad state: 102.1) than the predicted levels of 35 (good state) and 95 (bad state). These additional findings raise important questions: Why do buyers not use price announcements more often? Why do many buyers choose to make high price announcements? Where does the decrease in the shading rate in flexible contracts come from, given that buyers use informal agreements infrequently and often make high announcements if they use them?

In the Online Appendix we provide a separate analysis that investigates these questions in detail. Here we provide a short summary. Somewhat surprisingly, we find that buyers who decide not to make a price announcement do not suffer from a disadvantage. In fact, flexible contracts without price announcements lead to outcomes that are very similar to those under flexible contracts with low price announcements. However, this does not imply that the opportunity to make nonbinding price announcements is without consequence. On the contrary, a comparison of flexible contracts without price announcements in the IA and the BL shows that the opportunity to make informal agreements significantly changes outcomes even if buyers decide not to make use of the opportunity. In particular, in the good state flexible contracts without price announcements lead to significantly lower shading rates in the IA (0.11) than in the BL (0.23). It seems that in the presence of the opportunity to reach informal agreements the sellers interpret a buyer's choice not to attach a price announcement as a signal that the buyer does not plan to pay a high price.

Explaining why some buyers decide to make high price announcements turns out to be more difficult. Buyers who attach high price announcements to a flexible contract tend to end up with low profits for one of two reasons: either they stick to their announcement and lose money, because they pay a high price; or they violate their agreement by paying a low price, which seems to make sellers unhappy and leads to a substantial increase in the shading rate. There are two possible motives that might explain why some buyers nevertheless choose to make high price announcements.

First, there may be fair-minded buyers who simply prefer to pay a high price. This is consistent with the fact that some buyers stick to their high price announcements. Second, in the post-experimental questionnaire some buyers also wrote that they made high price offers to make the contract more attractive for sellers. This motive is somewhat misguided since contracts are attractive to sellers anyway and high price announcements which are later violated only increase the shading rate. See the Online Appendix for further details on these points.

4.2. Revision

In the RE treatment our focus is on the following two aspects. First, we investigate whether contracts remain reference points even when revision is feasible. This is important, because ex-ante contracts may no longer provide reference points once the trading parties know that it is feasible to change contracts ex post. Second, we examine how buyer-induced revisions affect sellers' performance and how this depends on the circumstances under which revision has been initiated. This analysis sheds light on the behavioral consequences of revision, an aspect that has so far received little attention in the literature.

To investigate whether contracts remain reference points in the presence of revision opportunities it is most informative to compare outcomes of rigid contracts in the good state of the world across treatments. If contracts no longer provide reference points one would expect to either observe price revisions or more shading in the RE than in the BL. The intuition behind this hypothesis is that in the RE sellers might hope that buyers use their revision power to increase the price. If a buyer decides to stick to the competitively determined contract (which typically implies that the seller gets a low price), the seller may be aggrieved and shade.

A comparison of the left-hand side of Table 2 (results of the BL) with its right-hand side (results of the RE) reveals that this hypothesis is not supported by the data. The competitive auction outcomes for rigid contracts are slightly lower in the RE (38.4) than in the BL (40.5), but the difference is not significant.[22] The vast majority of buyers who choose a rigid contract do not revise the contract when the good state of the world is realized (the relative frequency of revision amounts to 0.09). As a consequence, the sellers in rigid contracts in the RE face slightly lower prices (38.7) than the sellers in the BL treatment (40.7), but this difference is again not significant.[23] However, the most important result is that the shading rate is not higher in the RE (0.06) than in the BL (0.06).[24] Figure 3 provides a graphical representation of these findings. The fact that the same prices do not trigger more shading in the RE than in the BL suggests that the presence of revision opportunities does not induce sellers to hope for high prices outside the boundaries of the ex-ante contract. Even though sellers know that

22. Ranksum test (rigid contracts, RE versus BL): auction outcome: $p = 0.176$.

23. Ranksum test (rigid contracts, RE versus BL): price: $p = 0.157$.

24. Ranksum test (rigid contracts, RE versus BL): shading: $p = 0.829$.

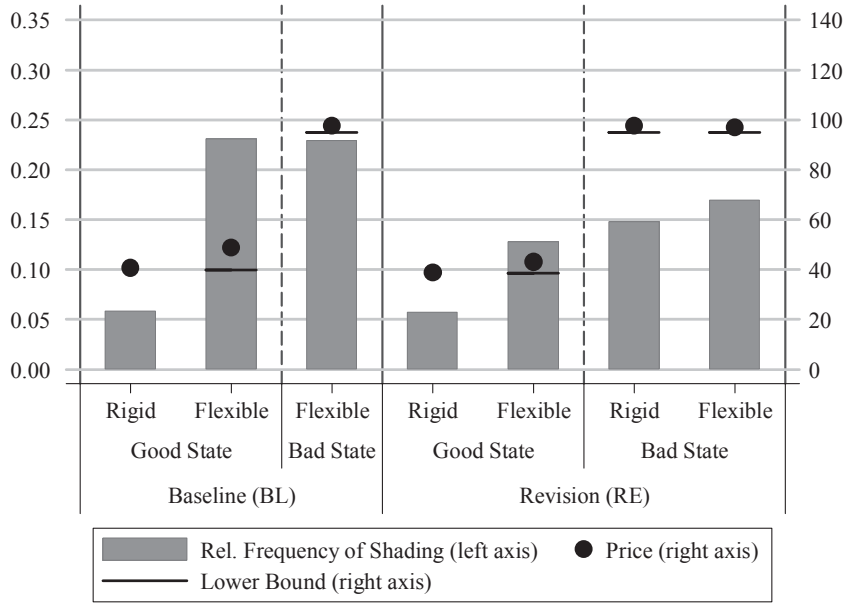

FIGURE 3. Baseline versus revision: prices and shading rates. The figure is based on data from the BL and RE. In the RE the data for the good state include nonrevised and revised contracts, the data for the bad state include revised rigid contracts (nonrevised rigid contracts do not allow for trade) and flexible contracts. Bars represent relative shading frequencies (left axis). Lines display the average lower bound of the price range, and dots represent average actual prices (right axis).

buyers could revise the contract and increase the price at no cost, sellers do not seem to expect buyers to do that. This suggests that contracts remain reference points even in the presence of revision opportunities.

Further support for this finding comes from column (1) of Table 4. We use observations from nonrevised rigid contracts in the RE and the BL in which the good state of the world has been realized and regress an indicator variable for shading on an indicator variable for the RE, the price, and an interaction term of the two. No coefficient in the regression is significant. This suggests that the availability of revision has not increased sellers' proneness to engage in shading in nonrevised rigid contracts even if the ex-ante agreed-upon prices are very low (the "RE" dummy measures the effect at the competitive price level of 35).

For the examination of the effects of revision on sellers' ex-post behavior we need to distinguish between three types of revision: (i) mutually beneficial revisions, (ii) opportunistic revisions, and (iii) altruistic revisions.

Mutually beneficial revisions occur when a buyer in a rigid contract initiates revision to increase the price after the bad state has been realized. Because trade would be infeasible without revision, both trading parties benefit (at least weakly) from this revision. In Section 3.2 we hypothesize that such revisions simply turn a rigid contract into a flexible one, because in the bad state a revised rigid contract and a flexible

TABLE 4. Shading in nonrevised and revised contracts in the RE.

Dependent variable	Shading	Shading	Shading	Shading
Treatment	BL & RE	RE	RE	RE
State	Good	Bad	Good	Good
	(1)	(2)	(3)	(4)
RE	0.001			
	(0.019)			
Price increment	0.000	−0.013	−0.001	0.000
	(0.001)	(0.007)	(0.001)	(0.001)
RE x Price inc.	−0.000			
	(0.002)			
MB revision		−0.010		
		(0.061)		
MB Rev. x Price inc.		−0.001		
		(0.008)		
Flexible contract			0.062*	
			−0.028	
OPP revision			0.483***	
			(0.130)	
Flex contr. x OPP rev.			−0.032	
			(0.185)	
ALT revision				−0.012
				(0.041)
Constant	0.056***	0.194***	0.042***	0.039***
	(0.015)	(0.059)	(0.009)	(0.009)
Observations	1295	298	1272	552
Clusters (Match. groups)	23	11	11	11
R^2	0.000	0.021	0.136	0.000

Notes: Column (1) uses data from nonrevised rigid contracts in the BL and the RE in which the good state has been realized. Column (2) uses data from revised rigid contracts and flexible contracts in the RE in which the bad state of the world has been realized. Columns (3) and (4) use data from rigid contracts in the RE in which the good state has been realized. In column (3) observations with altruistic revisions are excluded. In column (4) observations with opportunistic revisions are excluded. "RE" is an indicator variable which is unity if the observation comes from the RE. "Price increment" is equal to price minus 35 (good state) or price minus 95 (bad state). "MB revision" is an indicator variable which is unity if a rigid contract has been revised after the bad state of the world has been realized (mutually beneficial revision). "Flexible contract" is an indicator variable which is unity if the contract is of the flexible type and zero otherwise. "OPP revision" is an indicator variable which is unity if a rigid or a flexible contract has been revised in the good state of the world and the buyer lowered the price to a level below the auction outcome (opportunistic revision). "ALT revision" is an indicator variable which is unity if a contract has been revised in the good state of the world and the buyer increased the price to a level above the auction outcome (altruistic revision). All columns report coefficients of OLS estimations (linear probability model). Since observations within matching groups may be dependent all reported standard errors are adjusted for clustering at the matching group level.
***Significant at 1%; *significant at 10%.

contract allow for the exact same price choices. Table 2 and Figure 3 lend support to this hypothesis. Buyers in the RE frequently choose to revise rigid contracts in the bad state (the revision rate is 0.80). A comparison of outcomes in these revised rigid contracts and in flexible contracts reveals that both the price levels and the shading rates

are very similar (revised rigid contracts: price = 97.6, shading rate = 0.15; flexible contracts: price = 96.9, shading rate = 0.17).[25]

Column (2) of Table 4 shows the corresponding regression results. We use observations from contracts in the RE in which the bad state has been realized and compare revised rigid contracts with flexible contracts. To this end we regress an indicator variable for shading on an indicator variable for revised rigid contracts, the price, and an interaction of the two. The fact that neither the dummy for mutually beneficial revised rigid contracts nor the interaction term are significant confirms that outcomes of revised rigid contracts and flexible contracts are very similar in the bad state of the world in the RE.

Opportunistic revisions occur in the good state if a buyer decides to revise the price down to a level which is either below the ex-ante agreed-upon price (rigid contract) or the lower bound of the price range (flexible contract). Table 2 reveals that in the good state contract revisions are rare. Rigid contracts are revised in 9% of the cases. A closer analysis reveals that more than half of these revisions are altruistic revisions in which the buyer voluntarily increases the price. Opportunistic revisions therefore occur in only 4% (rigid contract) and 6% (flexible contract) of the cases.

Figure 4 displays the outcomes of nonrevised and revised contracts in the RE. The figure reveals that opportunistic revisions in which buyers simply attempt to grab a bigger share of the gains from trade lead to very high shading rates (rigid contracts: 0.52 / flexible contracts: 0.55). The corresponding regression analysis is found in column (3) of Table 4. We use observations from contracts in the RE in which the good state has been realized and regress an indicator variable for shading on indicator variables for rigid contracts and opportunistic revision and an interaction term of the two.[26] We also control for the price level. The regression shows that opportunistic revisions increase the frequency of shading massively (0.483, $p = 0.004$). The insignificant interaction term indicates that the large effect is present in rigid and flexible contracts (-0.032, $p = 0.866$). The finding that opportunistic revisions trigger strongly negative responses is important, because most existing incomplete contracting models abstract from the possibility that revision may lead to costly conflicts. Our results suggest that future theories need to take into account these psychological aspects of revision.

Finally, we also observe a few "altruistic" revisions in rigid contracts. "Altruistic" revisions are cases in which the buyer initiated revision to increase the price although the good state had been realized and trade would have been possible under the original contract. Column (4) of Table 4 investigates the consequences of altruistic revisions on sellers' shading behavior. Using rigid contracts in the good state in the RE, we regress an indicator variable for shading on an indicator variable for altruistic revision

25. Signrank tests (RE, revised rigid contracts versus flexible contracts): price: $p = 0.859$; shading: $p = 0.561$.

26. To provide a clean comparison between nonrevised rigid contracts and opportunistically revised contracts we exclude rigid contracts in which the price has been altruistically revised from the regression.

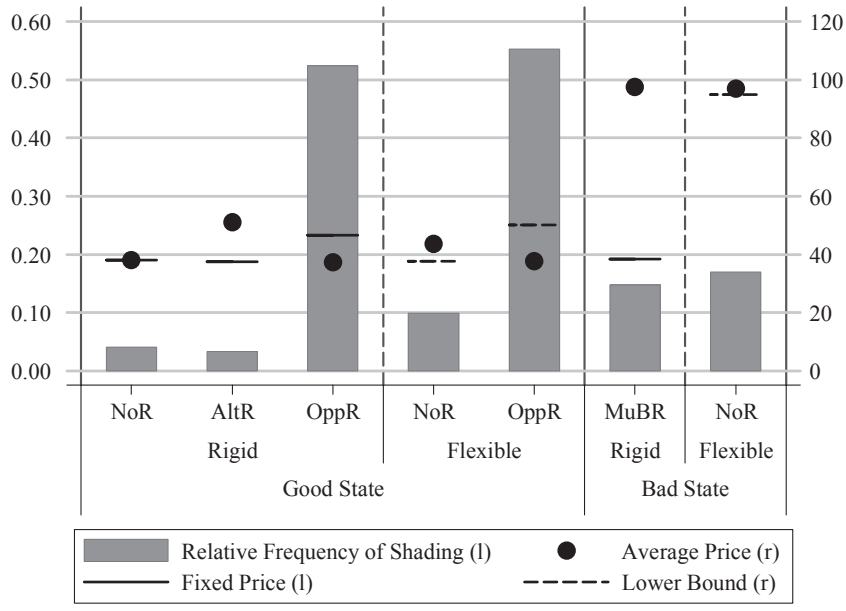

FIGURE 4. Revision: shading behavior in rigid and flexible contracts. The figure is based on data from the RE. "NoR" stands for contracts that have not been revised. "AltR" represents cases of "altruistic" revisions, i.e., rigid contracts in which the buyer has increased the price although the good state has been realized. "OppR" contains contracts in which the buyer has initiated an opportunistic revision to lower the price below the fixed price in rigid contracts or below the lower bound of the price range in flexible contracts. "MuBR" stands for mutually beneficial revisions. These are rigid contracts in which the buyer has increased the price to make trade feasible after the bad state has been realized. The bars represent relative shading frequencies (left axis). The solid lines display the average fixed price in rigid contracts, dashed lines show average lower bounds in flexible contracts and dots represent average actual prices (right axis).

and control for the price level.[27] Not surprisingly, the regression reveals that altruistic revisions do not trigger more shading (see also the graphical representation in Figure 4).

To summarize, the RE yields three important insights:

(i) Contracts continue to provide reference points even if revision is feasible. Although prices are typically low, rigid contracts are rarely revised when the good state is realized. Sellers seem to accept this behavior and do not shade more than in the BL.

(ii) Mutually beneficial revisions seem to turn rigid contracts into flexible ones. Prices and shading rates in revised rigid contracts and flexible contracts are similar in the bad state of the world.

(iii) Opportunistic revisions are perceived as hostile and trigger very high shading rates.

27. We exclude observations in which the price has been opportunistically revised.

In addition to these three main results, Table 2 also reveals another interesting finding. While the possibility of revising the contract has almost no effect on outcomes in nonrevised rigid contracts (relative to the BL), outcomes in flexible contracts seem to be substantially different. Both prices and shading rates are much lower in the RE than in the BL. The differences are most pronounced when the good state is realized, but they are also present in the bad state of the world.[28] These findings seem to suggest that sellers evaluate outcomes in flexible contracts differently depending on whether or not the buyer has the opportunity to revise the contract.

We explore this point in more detail in the Online Appendix. In short, our analysis suggests that the availability of revision affects sellers' perception of outcomes in nonrevised flexible contracts. Sellers seem to be more willing to accept low prices in flexible contracts if they know that the buyer has the opportunity to revise the contract. A possible explanation for this effect is that the presence of the revision opportunity shifts the focus partly away from prices to the revision decision itself. That is, there may be sellers who find all outcomes acceptable as long as the buyer does not revise the contract in opportunistic ways (see the Online Appendix for more details).

5. Related Literature

While the field of behavioral contract theory has only come to life in recent years, the literature is growing rapidly. On top of our own previous work on which this paper builds (discussed in the Introduction), there are several other papers that are directly related. Most relevant with regard to our informal agreement treatment is a recent study by Brandts et al. (2013) who (independently from us) also investigate how communication affects the frequency and effectiveness of flexible and rigid contracts. In stark contrast to our experiment they use a setup without ex-ante competition. Hart and Moore (2008) argue that ex-ante competition may be crucial to the extent that it is the objectivity provided by the market that turns the contract into a salient reference point. Interestingly, Brandts, Charness, and Ellman (2013) find that without communication rigid contracts are superior to flexible contracts even in the absence of ex-ante competition. This result suggests that contractual reference points may be more broadly applicable than hypothesized before. However, although the use of different parameters and procedures does not allow for a direct comparison of effect sizes, it seems that the trade-off between rigidity and flexibility that buyers face is less pronounced in their no-communication condition than in our baseline treatment where competition is present.

28. The comparison of outcomes of flexible contracts in Table 2 is not entirely clean, because the data in the RE also include flexible contracts which have been opportunistically revised. However, as our analysis has demonstrated, opportunistic revisions tend to increase the shading rate massively. Thus, if anything, the inclusion of opportunistically revised contracts biases the shading rate in the good state of the world upwardly. However, as opportunistic revisions are rare, this impact is negligible (we report the clean numbers in our more extensive discussion of this point in the Online Appendix).

In a second treatment, Brandts, Charness, and Ellman (2013) allow for the same restricted form of communication as we use in our informal agreement treatment (nonbinding announcements of state-contingent prices). They find that communication renders flexible contracts more profitable for buyers than rigid contracts. Also, in their setup the effects of this form of communication are rather weak, but because the advantage of rigid contracts is small in the absence of communication, the effect is sufficient to destroy the trade-off between rigidity and flexibility. In a third treatment, Brandts et al. introduce two-way, free-form communication over the full duration of the trading phase. They find that this more extensive form of communication substantially reduces conflicts in both types of contracts. As a consequence, flexible contracts clearly dominate rigid ones. The authors argue that free-form communication helps more than structured communication, because it allows the parties to establish a friendly relationship and helps to make promises more credible.

Whether free-form communication would also increase the attractiveness of flexible contracts in a competitive setup is an interesting and nontrivial question. In a setting in which a buyer faces multiple potential sellers, free-form communication becomes more burdensome as the buyer needs to communicate with many possible trading partners simultaneously before the contract is signed. In addition, competition may introduce incentives for strategic ex-ante communication which may undermine the credibility of promises. We think that this is a very interesting and important avenue for future research. Given our results and the findings of Brandts, Charness, and Ellman (2013), it is important that we understand more deeply what exactly determines the degree to which communication allows the trading parties to limit conflicts.

Regarding our revision condition, there are several related papers. Bartling and Schmidt (2014) conduct a laboratory experiment in which a buyer and a seller can trade a widget but have to agree on the terms of trade before knowing the optimal time of delivery. If it turns out that there is a better delivery time than the one the parties have agreed to, the buyer can ask for a change in the delivery time. In response, the seller can ask for a markup on the original price. To investigate the role of ex-ante contracts for the revision phase, they compare this main condition (the contract treatment) with a condition in which the same threat points which endogenously emerge in the main treatment are exogenously imposed on the trading parties (the no-contract condition). They find that sellers ask for much lower markups and buyers are more likely to reject given markups in the contract condition than in the no-contract condition. These findings suggest that the initial contract serves as a reference point. Moreover, these results are in line with our finding that revisions are quite unproblematic as long as they are seen as fair (which is typically the case if a seller asks for a low markup in response to a buyer's request for a change in the delivery time), but lead to counterproductive conflicts if one of the parties feels that the other party has behaved opportunistically (e.g., if the seller asks for a high markup).

Hoppe and Schmitz (2011) report a related result. In their experiment participants can use option contracts to avoid hold-up. Standard contract theory predicts that (some) option contracts are helpful only if renegotiation is not feasible (otherwise there is an incentive for a party not to exercise the option and renegotiate). However, the results of

the experiment reveal that option contracts also improve performance if renegotiation is feasible. It seems that buyers often refrain from opportunistic price renegotiation, because they anticipate that many sellers feel entitled to the option price stipulated in the contract and would refuse to trade at a lower price (even if refusal is costly to the seller). These results further support the view that many people perceive opportunistic renegotiations or revisions as hostile and are willing to punish such activities. Iyer and Schoar (2012) observe similar behavioral patterns in a field experiment in which they confront wholesalers of custom-made pens with hold-up threats. Many of the wholesalers are not willing to accept lower prices even if this implies that they lose a valuable trade.

The fact that similar behavioral patterns emerge in a broad set of different situations suggests that future theories of incomplete contracting would be well advised to take the behavioral effects of revision into account.

6. Conclusions

We show experimentally that the central behavioral mechanism underlying the concept of contractual reference points is robust to informal agreements and ex-post revision. Our experiment provides further evidence in line with the idea that contractual reference points importantly shape performance in trading relationships governed by incomplete contracts.

Our informal agreement treatment reveals that the possibility of making nonbinding price announcement lowers the shading rate somewhat, but the effect is rather weak, so that the trade-off between contractual rigidity and flexibility remains. Hart and Moore (2008) point out that the idea of managing reference points through informal agreements has some force, but they argue that asymmetric information in combination with self-serving biases may limit the impact of such agreements considerably. Our results illustrate that (at least in certain environments) asymmetric information is not necessary for the effects of informal agreements to be weak. In our setup both parties are fully informed about the cost level of the seller. In addition, our setup is very simple, in the sense that there is no uncertainty about buyer values and only two possible realizations of seller costs. This makes state-contingent informal agreements straightforward, because it is sufficient to determine a price for each of the two cost levels. We think it is therefore fair to say that we implement an environment in which informal agreements have a very good chance to be effective. In this sense the fact that we find a weak impact of informal agreements is a strong result.

Our revision treatment yields three important insights. First, our results indicate that contracts continue to be reference points even if revision is feasible. Sellers do not seem to hope for outcomes outside the contract even if the competitively negotiated terms are not very advantageous to them. Second, in line with Hart (2009), we find that opportunistic revision in itself has important behavioral effects. If buyers revise a contract for opportunistic reasons, sellers respond with a very high shading rate. Finally, we also identify circumstances in which a rigid contract that is revised can achieve the

benefits of flexibility without incurring its costs. In particular, it is sometimes better for parties to write a simple (rigid) contract and then revise it ex post if needed, rather than to anticipate and include future contingencies in a (flexible) contract from the outset.

It is worth (re-)emphasizing some limitations of our work. In our informal agreements treatment we have restricted attention to particular types of communication and we have supposed completely symmetric information between the parties. In our revision treatment we have studied revision in situations that are very simple, and where the temptation to revise opportunistically is low. Investigating the robustness of our results in richer environments is an important topic for future research.

References

Bartling, Bjoern and Klaus Schmidt (2014). "Reference Points, Social Norms, and Fairness in Contract Renegotiations." *Journal of the European Economic Association,* 12, forthcoming.

Brandts, Jordi, Gary Charness, and Matthew Ellman (2013). "Let's Talk: How Communication Affects Contract Design." Working paper, University of California, Santa Barbara.

Bolton, Gary E. and Axel Ockenfels (2000). "ERC—A Theory of Equity, Reciprocity and Competition." *American Economic Review*, 90, 166–193.

Camerer, Colin (2003). "Behavioral Game Theory: Experiments in Strategic Interaction." Roundtable Series in Behavioral Economics, Princeton University Press.

Charness, Gary and Matthew Rabin (2002). "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics*, 117, 817–869.

Dufwenberg, Martin and Georg Kirchsteiger (2004). "A Theory of Sequential Reciprocity." *Games and Economic Behavior*, 47, 268–298.

Falk, Armin and Urs Fischbacher (2006). "A Theory of Reciprocity." *Games and Economic Behavior*, 54, 293–315.

Fehr, Ernst and Simon Gaechter (2000). "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, 90, 980–994.

Fehr, Ernst, Lorenz Goette, and Christian Zehnder (2009). "A Behavioral Account of the Labor Market: The Role of Fairness Concerns." *Annual Review of Economics*, 355–384.

Fehr, Ernst, Oliver Hart, and Christian Zehnder (2009). "Contracts, Reference Points, and Competition—Behavioral Effects of the Fundamental Transformation." *Journal of the European Economic Association*, 7, 561–572.

Fehr, Ernst, Oliver Hart, and Christian Zehnder (2011). "Contracts as Reference Points—Experimental Evidence." *American Economic Review*, 101, 493–525.

Fehr, Ernst and Klaus Schmidt (1999). "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics*, 114, 817–868.

Fischbacher Urs (2007). "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, 10, 171–178.

Greiner, Ben (2004). "An Online Recruiting System for Economic Experiments." In *Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht 63*, edited by Kurt Kremer and Volker Macho. Ges. für Wiss. Datenverarbeitung, Göttingen, pp. 79–93.

Güth, Werner, Rolf Schmittberger, and Bernd Schwarze (1982). "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization*, 3, 367–388.

Hart, Oliver (2009). "Hold-Up, Asset Ownership, and Reference Points." *Quarterly Journal of Economics*, 124, 267–300.

Hart, Oliver and Bengt R. Holmstrom (2010). "A Theory of Firm Scope." *Quarterly Journal of Economics*, 125, 483–512.

Hart, Oliver and John Moore (2008). "Contracts as Reference Points." *Quarterly Journal of Economics*, 123, 1–48.
Hoppe, Eva and Patrick Schmitz (2011). "Can Contracts Solve the Hold-Up Problem? Experimental Evidence." *Games and Economic Behavior*, 73, 186–199.
Iyer, Rajkamal and Antoinette Schoar (2012). "The Importance of Holdup in Contract Negotiations: Evidence from an Audit Study." Working paper, MIT.
Klein, Benjamin (1996). "Why Hold-Ups Occur: The Self-Enforcing Range of Contractual Relationships." *Economic Inquiry*, 34, 444–463.
Maskin, Eric and John Moore (1999). "Implementation and Renegotiation." *Review of Economic Studies*, 66, 39–56.
Rabin, Matthew (1993). "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, 83, 1281–1302.
Rajan, Raghuram (2012). "The Corporation in Finance." *Journal of Finance*, 67, 1173–1217.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Online Appendix A: Additional Results
Online Appendix B: Replication of Results with 2013 Data Only