

The Importance of Hold-up in Contracting: Evidence from a Field Experiment

Rajkamal Iyer and Antoinette Schoar^{*}

Abstract:

This paper explores how the relationship specificity of an investment affects the ex-ante structure of contracts and the ex-post resolution of an ensuing hold-up problem. We set up a field experiment in the wholesale market for pens in India where we sent trained auditors as entrepreneurs to procure large orders of pens from wholesale dealers. We vary the specificity of the order by either buying generic or custom-printed lots of pens. We find clear support for the importance of incomplete contracting models. The greater threat of hold-up in the case of printed pens leads wholesalers to demand a significantly higher upfront payment compared to generic pens. Similarly, when faced with actual hold-up, wholesalers are more willing ex-post to renegotiate an order and accept lower prices in the case of printed pens. But while the signs of the contract terms corroborate the intuitions of incomplete contracting models, they do not explain the magnitudes we observe. First, the level of the upfront payment is so low that it does not cover the procurement costs of the wholesaler. And second, we find that wholesalers often do not use their bargaining power even when a shopper holds them up ex-post. In fact we observe that in 60% of the generic deals they return the upfront payment when the shopper tries to hold them up. We also find that there is a lot of variation in the contract structure that is not explained by the relationship specificity of the investment.

^{*} Iyer: MIT Sloan School of Management, email: riyer@mit.edu; Schoar: MIT Sloan School of Management, CEPR and NBER, email: aschoar@mit.edu. We thank Bade Kucukoglu, Janina Matuszeski and especially Sandhya Kumar for outstanding research assistance. We thank Robert Gibbons, Oliver Hart, Sendhil Mullainathan, Gordon Phillips, Morten Sorenson, Chris Woodruff and Luigi Zingales for many helpful comments. The Institute for Financial Markets Research in Chennai, India provided financial support. All errors are our own.

1. Introduction

Incomplete contracting theories have been a central building block for models in finance and economics for almost three decades, starting with Klein, Crawford and Alchian (1978) and Williamson (1979) and further developed by Grossman and Hart (1986) and Hart and Moore (1990). These theories build on the idea that when state contingencies cannot be perfectly observed, described or enforced, contracting parties are often left to settle contractual issues via ex-post (re)negotiation. This in turn can lead to ex-post hold-up of the party that has made greater relationship specific investments. In a world of forward looking actors the threat of ex-post expropriation can preclude parties from making relationship specific investments or exerting effort upfront.

Despite the widespread use of these theories in organizational economics and corporate finance very few empirical studies directly test the extent to which hold-up plays a role in explaining how actual contracts are structured ex-ante and whether parties indeed resort to hold-up when the opportunity arises.¹ Similarly, we only have very little evidence on how contracts are renegotiated in case one party does resort to hold-up and whether the outcome is affected by ex-ante contract structures. The reason is that the hurdles for convincingly testing these theories are high since it requires exogenous variation in contractual structures in addition to collecting detailed data on contracts and outcomes.

To address these questions we conduct a randomized field experiment in the Southern Indian city of Chennai that explicitly analyzes the importance of hold-up when parties have to make relationship specific investments and contracts are incomplete. We use a novel audit study methodology where we send trained auditors to negotiate real purchase agreements in the wholesale market for pens and execute the purchase. To exogenously vary the level of relationship specific investment the wholesalers have to make, we randomly place half the orders for generic, plain pens and the other half for printed pens that have no resale value outside the relationship with our auditors. The possibility for hold-up arises ex-post since shoppers can renege on the agreed upon price as the contracts are not enforceable in this market. Thus a basic prediction of the incomplete contracts model is that the demanded upfront payment should be higher for the printed than the generic pens.

We analyze how the potential for hold-up affects the ex-ante structure of contracts and the renegotiation ex-post. Our experiment produces a number of interesting insights. On the one hand we find clear support for the importance of relationship specific investments in contracting. The greater threat of hold-up in the case of printed pen orders leads wholesalers to demand a significantly higher fraction of the payment upfront than in the case of generic pens (36% versus 11%). Yet on average there is no difference in the pricing of the generic versus the printed pens, which makes sense in the context of an incomplete contracts model since the initial contract price can be renegotiated ex-post.

¹ While there is some limited evidence to corroborate that asset redeployability affects the sale value and price of assets, see for example Pulvino (1998), Benmelech and Bergman (2008), there is no reliable evidence about the extent of hold-up in contracting situations where parties make relationships specific investments.

Similarly, in line with the intuition of the incomplete contracts literature, we find that ex-post, when faced with actual hold-up, wholesalers are more willing to renegotiate the order in the case of printed pens and are prepared to accept a price below the cost price of the pens. Neither of which happens in the case of generic pens. In addition, the magnitude of discount that is offered by a wholesaler in renegotiation is lower if the upfront payment he had received before is larger, which suggests that wholesalers are less willing to renegotiate if they have more bargaining power when the shopper already paid a large fraction of the money upfront. These findings support the hypothesis that the ex-ante contract structure is set up to mitigate future hold-up risk due to specificity of the investment.

But on the other hand, we also find that incomplete contracts interact with market forces in a more complex way than static models of hold-up predict. While these theories get the signs of the contract terms right they do not explain the magnitudes we observe. First, the level of the upfront payment which wholesalers demand for printed pens is very low (on average it is only 40% of the purchase price). This does not cover the cost of procuring the pens from the manufacturer, which leaves the wholesaler exposed to the risk of hold-up and ultimately puts the survival of the business at risk. Now wholesalers could still break even if they themselves expected to hold-up customers who come to pick up their orders. However, in all our visits, there was not a single situation where a wholesaler tried to renegotiate the price. Therefore wholesalers on average must expect to see few incidences of hold-up in this market even though customers would always benefit from doing so.

By itself, low prevalence of hold-up would not preclude wholesalers from requesting a level of upfront payment that is high enough to cover their procurement costs. However competition among wholesalers forces them to minimize not only the price but also the level of upfront payment they charge. If there is a distribution of clients with different propensity to renege on a contract, a wholesaler who is better at estimating the type of a client can offer more favorable price and contract structure. This means wholesalers have to take into consideration the type distribution of clients and set the right level of upfront.

A second major deviation from static hold-up models is that wholesalers often do not seem to use their bargaining power even when a shopper tries to hold them up. In the case of generic pen orders we find that most wholesalers decline to renegotiate the order and state that they would rather sell the pens in the market. But interestingly, in 60% of the failed renegotiations, wholesalers offer to return the upfront payment that the shopper had already made. The same does not happen in the case of printed pens. This is very surprising since wholesalers would always be better off to keep the money. It suggests that wholesalers' trade off the incentives to hold-up their customers with either reputational concerns or prevailing social norms.

Finally, we also show that there is a large amount of variation in contract structure and even more in the pricing of the orders that is not explained by assets specificity. Instead, a much larger fraction of variation in the data is explained by wholesaler and shopper fixed effects. We find an adjusted R-squared of 10% in the former case versus over 70% when we include the wholesaler and shopper fixed effects. This suggests that a major fraction

of the behavior of the participants in this market is explained by forces outside the incomplete contracts framework.

We chose the wholesale market for pens in India as the venue for our audit study for a number of reasons: First, pens are an item that is relatively standardized in its quality but can easily be changed in the specificity of the order by printing something on it. Second, India is an ideal test case for incomplete contracting situations since the Indian court system is notoriously difficult to access and therefore any breach in a contract in this market is difficult to enforce. Third, this market has a large set of wholesale stores that are comparable in size and customer base which allows us to conduct a valid experiment.

In the first step of the experiment we randomly assign 46 auditors to pose as small business men (“shoppers” hereafter) and visit 107 wholesale stores for a total of 494 visits. The shopper negotiates a bulk order of pens with a wholesaler. The average order size is around 600 pens. Thus in total across 494 visits we purchase around 300,000 pens. The order consists either of generic pens that are easy to resell in the market (“plain pens”) or customized pen orders where we ask the wholesaler to print a specific logo on the pens that cannot be removed once it is printed (“printed pens”). Wholesalers have to make a relationship-specific investment by ordering the pens from a large distributor and paying for the order upfront (and in the case of printed pens they also have to bear the printing cost). Printed pens require a greater relationship-specific investment from the wholesaler, since once the pens are printed they cannot be used in any other sale. Therefore, ex-post the shopper has more bargaining power to renegotiate the price of the order. But it is a two-sided hold-up situation since the wholesaler could also behave opportunistically and not deliver on time or try to charge a higher price at the time of delivery.² We believe that opportunistic behavior by the wholesalers in general is less likely since it would hurt their reputation.

A straightforward model of incomplete contracts would suggest that the amount of upfront payment demanded should be higher for the printed pens, since shoppers will attempt to hold-up the wholesaler ex-post. In fact, if wholesalers cannot hold-up shoppers ex-post, e.g. due to reputation reasons, the upfront payment should cover at a minimum the procurement price of the pens. As discussed above, we find that the required demanded upfront payment is higher in the case of printed pens relative to plain pens, which suggests that wholesalers respond to the threat of being held up ex-post when making a relationship-specific investment. It is interesting to note that not only do the wholesalers start by asking for a significantly higher fraction of upfront payment in the case of printed versus plain pens (58% versus 19%). But they are also less willing to be bargained down in the case of printed pens than in the case of plain pens. In contrast, when we look at the prices that are charged for the pens, we do not find a significant difference between printed and plain pens once we deduct the actual cost of the printing from the price of printed pens.

² In fact if there was no concern about hold up from the side of the wholesaler, the optimal contract would be for the shopper to pay the full price upfront. But we rarely ever observe that wholesalers charge close to 100% upfront which suggests that shoppers perceive the two-sided hold up problem.

But the level of the upfront payments is relatively small on average. Even in the case of printed pens the final contracted payments are only 35% of the total sales price. Since the resale value of printed pens in the market is zero and the profit margin on the pens is less than 15%, the required upfront payments cover less than 50% of the costs.³ If wholesalers were expecting to be held-up each time they would not be able to break-even given the observed contracts, since they would lose money on each transaction by being bargained down to the level of the upfront payment. Therefore, wholesalers have to be operating with an assumption that there is a significant fraction of shoppers that will never hold them up.

In a second step, we then conduct an ex-post renegotiation of the contracts for 75 pen orders (half of these orders are for plain and half for printed pens): After placing the initial order with the wholesaler, the shopper calls up the wholesaler on the agreed delivery day to announce that he is unable to honor the contract in full and instead asks to renegotiate the contract to a lower price. We again find evidence that the specificity of the investment affects the bargaining position of the wholesaler. In the case of printed pens, wholesalers are willing to renegotiate the order at a significantly higher frequency than in the case of plain pens. In addition the discount that is offered in these cases is much higher for printed than plain pens. Plain pens are never renegotiated below their procurement price, however, in 30% of the cases the re-contracted price for printed pen orders is below their cost price. But we also find that in 35% of the cases, wholesalers are unwilling to renegotiate the contract even for printed pens or abruptly terminate the renegotiations. This is surprising given that the salvage value for these printed pens is close to zero. This is a deviation from the assumption that renegotiation leads to ex-post efficient outcomes. However, this finding could be explained by wholesalers being concerned about reputation costs incurred in the market.

Examining the impact of the ex-ante contract structure on the renegotiation, we find that the level of the upfront payment that was provided initially affects the willingness of the wholesaler to renegotiate the deal. The higher is the upfront payment, the lower the discount that the wholesaler is willing to accept, especially in the case of printed pens. This result suggests that wholesalers take into account their bargaining position and the outside value of the pens when deciding how to renegotiate a contract.

But as in the upfront contracting above we see that there is a significant subset of wholesalers who do not use their bargaining power to hold up the shopper when they could. In 15 out of 40 plain pen contracts the wholesaler in fact agrees to return the upfront payment to the shopper in full. This violates the key assumption that either party will renegotiate the contract in order to extract the full rents from the relationship. In this case, wholesalers should keep the upfront payment and not deliver the pens, since the shopper is breaking the contract. It is interesting that the wholesalers only offer to return the upfront payment in the case of plain pens, where the wholesaler does not have to incur a loss due the renegotiation. In fact wholesalers are more likely to return the upfront when it is a small fraction of the total price. One could argue that this behavior is more consistent with reputational constraints than social norms. The wholesalers seem to trade

³ We know the procurement costs of the wholesaler.

off the money lost by returning the upfront payment with the potential reputational gain and decide to be generous when the amount they lose is small. If they were compelled by social norms to return the money, we might have expected that they would not take the amount of upfront into consideration.⁴

The remainder of the paper proceeds as follows: Section 2 places our paper in the context of the existing literature, Section 3 describes the experimental setup and the randomization approach. Section 4 describes the data. Section 5 discusses the results of the upfront contracting results and Section 6 summarizes the results from the ex-post renegotiation. Finally Section 7 concludes.

2. Literature Review

Our study is related to a number of different strands of the literature. By empirically examining the importance of hold-up among contracting parties, our paper sheds light on a key driver of the existing theories of contract incompleteness. Klein, Crawford and Alchian (1978) and Williamson (1979, 1985) argue that contract incompleteness in the presence of investments in specific assets opens the door to hold-up. They propose vertical integration as a solution to deal with the problem of hold-up.⁵ Grossman and Hart (1986), and Hart and Moore (1990) develop the property rights theory of the firm where ownership of the asset helps reduce the hold-up problem and increase the incentive of parties to make non-contractible relationship-specific investments. Aghion and Bolton (1992) and Bolton and Scharfstein (1996) develop a theory of capital structure based on control rights allocation with the main ingredient being incompleteness of financial contracts. Our paper also contributes to the existing literature that explores the importance of asset liquidation values in contracting. Williamson (1988) and Shleifer and Vishny (1992) highlight the role of asset redeployability in the determination of debt capacity. On the empirical front, Benmelech and Bergman (2008) find that airlines renegotiate their lease obligation downwards when their financial position is poor and their fleet liquidation value is low. Similarly, Acharya et al. (2007) find that creditors of defaulted firms recover less if the industry is in distress and non-defaulting firms in the industry are illiquid. Furthermore, they find the effect to be pronounced if debt is collateralized by specific assets that are not easily redeployable.

A growing empirical literature tests the role of reputation and trust for contract completeness. Crocker and Reynolds (1993) investigate the procurement contracts used by the U.S military and find that higher reputation and complexity lead to drafting a more incomplete contract. Banerjee and Duflo (2000) show that contracts written between firms are associated with the reputation levels of the firms. McMillan and Woodruff (1999) find that inter-firm trade credit is more likely when the delivering firm trusts the

⁴ Of course, one could imagine a more complicated model of social norms or fairness where the level of generosity expected from the wholesaler should be commensurate with the loss on the other side, but that story seems much more complicated to motivate.

⁵ Joskow (1985) finds higher incidence of common ownership in the form of vertical integration among electricity generating plants that site next to coal mines. Joskow (1987) finds that ex-ante contracting parties enter into longer contracts when relationship specific investments are more important.

client. Kaplan and Stromberg (2002) find that a central feature of venture capital contracts is the allocation of control rights between venture capitalist and entrepreneur suggesting that contracts are inherently incomplete.

Similarly there is a large literature using lab experiments to test the role of fairness in negotiations. These lab experiments highlighted that in many bilateral bargaining set-ups a non-negligible fraction of participants do not care solely about their material payoffs, but have fairness considerations (Güth et al., 1990; Roth 1995; Camerer and Thaler 1995).⁶ We build on the findings from these controlled lab settings by testing bargaining behavior in real contracting situations.

Finally, our methodology draws on the existing literature of audit studies even though their context and the questions differ completely from the current study. These audits focus mainly on discrimination due to auditor characteristics such as minority or gender status. For some of the most prevalent studies see Ayres and Siegelman (1995), Newmark et al., (1996) or Bertrand and Mullainathan (2004). We expand the approach of audit studies by engaging in real purchase transactions to test contracting and renegotiation.

3. Description of Experimental Set-Up

3.1. Market Structure and Intervention

The field experiment was conducted in Chennai, a city in Tamil Nadu, a state in the South of India. Chennai is the largest city in Tamil Nadu with over 4.5 million inhabitants. For the purpose of the study, we hired auditors to purchase large orders of pens in bulk from wholesalers. To test the differences in the bargaining outcomes for goods of different degree of relationship specificity we chose a single industry, pen and stationary wholesalers, to conduct our transactions. These are wholesale traders in pens and stationary items, not small stationary shops. The choice of this industry was driven by a number of different factors: (a) We were looking for an industry with a large number of similar sized establishments in a given location. Pen and stationary wholesalers provided such an opportunity, since there are more than 100 wholesalers in the Chennai area. Moreover, there is a distinct cluster of wholesalers in a particular neighborhood, called Paris Market. The benefit of such an arrangement is that it minimizes the amount of firm-specific shocks. In addition, the wholesalers are supplied by a single large distributor at a pre-set price. This helps us obtain a good estimate of the procurement price (cost) of the pens for the wholesalers.⁷ (b) We wanted to ensure that the type of commodity used in our experiment is relatively standardized, which in turn facilitates comparisons of deal terms offered by different wholesalers. But at the same time we wanted to make sure that the type of good provides opportunities to include customized features which will allow us to vary the potential threat of hold-up between the shopper and the wholesaler. The pen industry offers a great opportunity since we will be able to customize pens by printing customer specific logos on the pens. (c) We wanted an industry where first time buyers could place large orders without a prior history of

⁶ See also Fehr and Schmidt (1999).

⁷ The procurement price of pens is Rs 3.8 and the average profit margin of a wholesaler is 15%.

interactions. The pen industry met this requirement as there are a good proportion of first-time buyers placing large orders.⁸ Finally, the pen industry provides us with a cost-effective product of entering into bulk deals.

To implement the actual audit study, we hired auditors who themselves are entrepreneurs so that they are familiar with the process of bargaining for supplies and services. We verified the entrepreneurs we hired are not affiliated with the pen industry to avoid any familiarity between the wholesalers and the shoppers. Instead we looked for entrepreneurs from similar types of businesses such as grocery store owners, small manufacturers, etc. Once the auditors were hired, they were given training to explain the set-up of the experiment, the details of the pen industry and their particular assignment. The auditors were paid a fixed fee per visit to the wholesaler that was above market rate. They were told that in case of any discrepancies from the script they would not be invited to do further visits. The auditors were given the information that they are part of a study to understand contract terms in the pen industry. However, auditors (hence forth referred to as shoppers) were not told what the expected outcome of the study was in order to avoid any “demand effects” in their behavior. We also provided shoppers with a specific identity such as the type of firm they run. Most of the business profiles that the shoppers were assigned were firms such as advertising companies, event management firms etc., which justify why the shopper is placing a bulk order for pens. Shoppers are given the name of the firm they operate, the name and logo of the client for whom they are placing the bulk order and a business card with the associated information to credibly signal to the wholesaler that they run a legitimate business.

We divide the pen purchase by the shoppers into two types. In half the cases the shoppers are assigned to buy plain pens and in the other half they are assigned to order printed pens. A printed pen has a firm name or message embossed on it. In the case of plain pens no changes are done to the pen and it is a generic commodity. We also ensured that the shoppers are ordering only the most common types of pens so that there are no differences in the resale possibilities for the plain pens. The distinction between the plain and the printed pen is that the printing on the pen cannot be easily removed and thus destroys its marketability (also it takes on average 3 days for the wholesaler to get the printing done).⁹ The resale value of printed pens is close to zero as there is no easily available secondary market to retail these pens. Furthermore, wholesalers are not in the business of retailing thousands of secondhand pens and their distribution network is not geared for this.

The shoppers are asked to place a bulk order of a pre-specified size which we randomized across shoppers. The order size varies between 500-750 pens in increments of 50 pens. This range-of-order size was determined after conducting a number of pilot interviews with wholesalers who confirmed that this is the modal order size in the wholesale

⁸ From our pilot interviews, we gathered that 75% of the customers in this market are first-time buyers and only 25% are repeat buyers.

⁹ It is common for customers (marketing companies, event management firms, etc.) in the market to place orders for either printed pens or plain pens. Thus placing an order for printed pens does not signal a different possibility of repeat interactions with wholesalers.

industry. The profits from the order size of 500-750 pens correspond to around 5% of a wholesaler's weekly profits. Our aim is to mimic a regular business transaction that is neither too large to draw suspicion nor too small to be negligible to the wholesalers.

We also provided the shoppers with a detailed script that specified a bargaining rule which they were asked to follow during purchases. The visit to each wholesaler can be summarized as follows: At first the shopper would enter the establishment and buy some items from the wholesaler. Generally these were 20-25 pens of commonly available brands. The shopper would pay for these pens with cash to establish trust between the wholesaler and himself. At the time of making the payment, the shopper would introduce himself to the wholesaler. After the introduction, the shopper would mention to the wholesaler that he is interested in placing a bulk order of between 500 to 750 pens and inquire about the rate. Once the wholesaler makes the first offer, the shopper would make a counter-offer that is equal to the wholesale procurement price of the pen + 0.10 Rs. The second and third offer from the shopper would be at 0.10 Rs. higher than the previous offer. The fourth offer would be 0.10 Rs. lower than the previous offer of the wholesaler. The bargaining ends at any point if the wholesaler agrees on the price or refuses to bargain anymore.

Once the wholesaler and shopper have agreed on the price, for printed pens, the shopper would inquire about the printing rate. After fixing the printing rate, the shopper would inquire about the delivery time. If the delivery time were less than a week, the shopper would agree or else would negotiate for delivery in a week. After finalizing the delivery time, the shopper would negotiate about the mode of payment at delivery. The shopper would ask first for credit at the time of delivery, then propose a post-dated check, then check payment and finally cash payment at delivery. After the payment terms were fixed, the shopper would ask if he could return the pens in case they are defective. Finally, the shopper would tell the wholesaler that he will come in a week to collect delivery, at which time the wholesaler would invariably ask for an advance.¹⁰ The buyer would then negotiate for the advance payment using the following rule: 10% of the total amount, then 25% and finally 10% lower than the wholesaler's final offer.¹¹ It is important to note that we ask buyers to complete the deal at the lowest price possible, but we did not ask them to terminate the negotiation in case a certain price is not achieved. This means our experiment does not allow movement along the extensive margin, where some deals might not be reached if the wholesaler insists on a high price.

The bargaining process for plain pens is very similar to the printed pens on all dimensions except for the delivery time. For the delivery time, the shopper tells the wholesaler that he would like to take delivery after a week even if the wholesaler has the stock ready earlier.¹² To detect if the shoppers deviate from the script and also to check their performance, one of the visits of the shopper was to a wholesaler where the wholesaler was our representative (the shoppers were never informed about this).

¹⁰ The structure of the negotiation was based on what is prevalent in the wholesale market for pens.

¹¹ For the visits where the contract was renegotiated, the buyers were given a maximum limit of 40% for the final advance payment.

¹² The shopper states that he has many other appointments lined up and it would be difficult to carry the order with him throughout the day so he would prefer to take delivery at a later date.

Furthermore, in some of the other visits, we also had our representatives (whom the shopper was not aware of) visit the wholesaler at the same time as the shopper and witness the bargaining process.

For the visits where the contract terms are renegotiated, the renegotiation is carried out using the following bargaining script: At the date when the delivery of pens is to be picked up, the shopper would call up the wholesaler and say that there is a problem. The shopper would tell the wholesaler that the company on whose behalf the shopper had placed the order does not want to take delivery. The shopper also conveys to the wholesaler that he has no immediate use for the pens. Finally the shopper tells the wholesaler that the situation is problematic for both of them, so the best he can offer is to take delivery of the pens if the wholesaler offers him a discount. At first the shopper asks the wholesaler for a 30% discount of the contracted price. The second offer is for 20% discount of the contracted price. The third offer is for 10% discount of the contracted price and final the offer is for 5% discount of the contracted price. The bargaining ends at any point if the wholesaler agrees on the discount or refuses to bargain any more. If the wholesaler agrees to give a discount, the shopper informs the wholesaler that he would come later to collect the delivery at the re-contracted price. In case the wholesaler refuses to give any discount, the shopper asks the wholesaler if he can get a refund of the advance that he has paid. In case the wholesaler refuses to refund the advance the negotiation is terminated. Note that the negotiation can also be terminated at any point by the wholesaler.¹³

Directly after each visit/renegotiation the shopper is asked to fill out a detailed exit survey that asks about the outcome of the negotiation. The shopper also goes back to the wholesaler to take delivery at the agreed-upon time and pays the outstanding part of the bill.

3.2. Methodology of Randomization

The randomization involved matching 46 shoppers to 107 wholesalers and determining the characteristics of each visit for a total of around 500 individual visits. Each shopper was assigned to visit 11 different wholesalers. Each wholesaler was visited between 3 and 6 times, with the majority of wholesalers being visited 5-6 times. Most importantly, the randomization imposed that half of the visits to each wholesaler were plain pen visits (shopper ordered plain pens) and the remaining were printed pen visits. To test how the specificity of the good interacts with other dimensions of a business transaction, we randomly assign variation in the type of the order that the shopper places, such as the type of pen ordered, and the number of pens ordered. We also tried to keep the script for each shopper as consistent as possible across the 11 visits they made to various wholesalers.

To achieve these goals of variation in visit characteristics, while maintaining a similar script across visits for the shoppers, the randomization was calculated in four main steps.

¹³ An hour after the renegotiation call, the shopper called up the wholesaler and informed him that the client has reinstated the order and he would pick up delivery at the initial contracted terms. Thus the final delivery was collected after making the payment in accordance with the initial contract terms

First, each wholesaler and shopper were randomly assigned a profile with their “intrinsic characteristics.” Each shopper’s ethnic group was, of course, the shopper’s actual ethnic group. The other shopper characteristics were assigned randomly to create variation in the type of transactions. The main dimensions of variation are (1) the shopper’s “company type” was assigned from among four categories (event manager, advertising agency, conference organizer and marketing company). The idea was to pick four different types of firms that are very common in India and justify frequent bulk orders on behalf of other companies. It was important to vary these profiles to avoid suspicion in the minds of the wholesalers in case of frequent interactions with people that have very similar profiles. Shoppers were given business cards with the name and the “fictional” company that they owned. (2) The pen type was assigned from two different pen brands that are of very similar quality and price. Again this dimension was included to create variation in the type of orders that our shoppers place. (3) The number of pens ordered, in one bulk order varied between 500 and 750 pens in increments of 50. This dimension also added to the variation in the type of orders. (4) The number of days for pens to be held is the time that the shoppers request the bulk order to be ready. This dimension was varied to create heterogeneity between shoppers. This dimension is only binding in the case of printed pens, since plain pens are usually available within a day. However, even for plain pens the delivery was collected at a later date even if the wholesaler had ready stock. (5) And finally, we assigned about half of the visits to be printed pen visits (shoppers buy printed pens and the other half to be plain pen visits (shoppers buy plain pens). For visits with printed pens the shoppers were given different logos each time for the wholesaler to print on the pens.

The randomization also restricted the assignment of shoppers based on the location of the wholesaler, which we will refer to as a “location group.” Wholesale establishments located near to one another (so that wholesalers can see who is visiting a neighboring wholesaler) were assigned the same location group number. Shoppers would not be assigned to other wholesalers in the same location group. The idea behind this constraint is that it might create awkward interactions for the shoppers if a wholesaler who was previously visited sees the same shopper go to a neighboring wholesaler.

In a second step, shoppers and wholesalers were randomly assigned to one another in a constrained manner. One wholesaler and one shopper were selected randomly from among the group of wholesalers and the group of shoppers. The randomization program then checked that the shopper had not been previously assigned to visit a different wholesaler in that same location group (to avoid the same shopper visiting neighboring wholesaler), and that the wholesaler did not have a previously assigned visit by a shopper of that same ethnicity or company type. If these conditions were met, then this shopper and wholesaler pair was declared a match and the shopper was assigned to visit that wholesaler. The information about the shopper and wholesaler’s assigned visits was updated to reflect the new match, and both shopper and wholesaler were returned to their respective common pools to be available for future random matches. Once a wholesaler received five visits it was removed from the pool of available wholesalers, while shoppers with 11 visits were removed from the pool of shoppers. The result of this randomization was that the 46 shoppers were each assigned to visit 11 wholesalers, and each wholesaler had a range of shopper types assigned to visit it.

The third step in the randomization was to assign whether the visit was for a plain or a printed pen. Out of the five visits to a wholesaler, two were randomly selected to be plain pen visits, two were randomly selected to be printed pen visits, and the remaining visit was randomly selected to be either a plain or a printed visit. By assigning each visit to be either a printed or plain visit, based on a random stratification at the level of the wholesalers, each shopper ended up with a number of plain visits and a number of printed visits, typically 4-7 visits of each type. Finally, each printed pen visit was randomly assigned a logo to be printed on the pen.

A fourth step involved adjusting a few parameters randomly to avoid potential detection by wholesalers. Since the shoppers were randomly assigned one pen type, one number of pens to be ordered, and one number of days for pick-up, by chance a few wholesalers had multiple shoppers with similar profiles. It was determined that having three or four shoppers request the same type of pen or the same number of days before the pick-up of the pens would not arouse suspicion. However, having three or four shoppers request the same number of pens could arouse suspicion. Accordingly, for the wholesalers at which three or four shoppers were set to ask for the same number of pens, we randomly selected one or two shoppers to request 650 pens for their visits to this wholesaler only.

Throughout this randomization, all characteristics were assigned randomly, in either an unconstrained, constrained or stratified manner. The only aspect of the randomization that was not strictly randomly assigned was the relative timing of the visits, although there was still a great deal of randomly induced variation in this variable. For the most part, visits to different wholesalers by the same shopper were made in a random order, based on the randomly assigned characteristics of the visits. One additional step to the randomization was that some of the initial wholesalers we had selected for the study stopped selling pens or shut down (even before a single visit was made to the wholesaler). Any visit that was originally scheduled to a wholesaler no longer selling pens or no longer in business was replaced by a visit to an existing or a new wholesaler. Ideally each “bad” wholesaler would be replaced by one new wholesaler and all remaining visits that were scheduled for the “bad” wholesaler would go to the new replacement wholesaler. In practice, there were not enough new wholesalers to take the place of the “bad” wholesalers. With X new wholesalers available, we randomly selected X of the “bad” wholesalers to be replaced by a randomly selected new wholesaler. For the remaining “bad” wholesalers, for each visit, an existing wholesaler was randomly selected to have the visit go to that wholesaler, meaning this replacement wholesaler then would have 6 visits in total.

4. Data Description

In total we conducted 494 audit visits to 107 different wholesalers by 46 different auditors, also called shoppers. The summary statistics in Panel A of Table 1 shows that the average wholesaler was visited 4.6 times with a minimum of 2 and maximum of 6 visits; the modal wholesaler received 5 separate visits by different buyers. The modal shopper completed 11 visits with a minimum of 6 and a maximum of 13 visits. We randomly assigned the order sizes of the purchase to the shoppers; these are evenly distributed between 500-750 with the most common lots being 550, 600, and 700. Panel 1 shows the average size of an order is approximately 615 pens. The average order size for

printed pens is 619 and for plain pens is 616.¹⁴ To ensure that our random assignment was conducted successfully, we also verified that other wholesaler and visit characteristics such as wholesaler location, time to delivery and ethnic background of the wholesalers and shoppers does not vary significantly between the printed versus plain orders. Panel B of Table 1 shows the statistics for the ex-post renegotiation of orders. We randomly picked 75 of the wholesalers and added one final visit where we renegotiated the order afterwards. We engaged 15 shoppers for that exercise. However, each of these 75 wholesalers only received 1 renegotiation visit, since we did not want to place an undue burden on the wholesaler.

In Table 2 we now report the average statistics for the different dimensions of the bargaining outcomes. The two most important dimensions of the contract are the price of the pens and the fraction of upfront payment that has to be made at the time of placing the order. In the case of printed pens there is an additional dimension, which is the cost of the printing. The descriptive statistics in Table 2 map out the negotiation process of the wholesaler visits. From Panel A we see that the fraction of printed to plain pens is roughly balanced, with 240 visits for printed pens and 254 visits for plain pens. The difference in the sample size stems from the fact that some wholesalers where we sent shoppers to ask for printed pens were not able to do the printing and we had to drop these visits. The initial price that is offered by the wholesaler on average is Rs 5.3 for the printed pens and Rs 4.9 for the plain pens. This difference is not surprising since the printing costs are around Rs 0.4. The final offer after bargaining is Rs 4.9 for the printed pens and Rs 4.5 for the plain. So on average the buyers receive about a 10% price reduction after bargaining. Again, the spread between the two groups of pens stays constant, verifying that the different in price is mainly a reflection of the printing costs. But there is wide variance around the price, with a minimum upfront price for printed pens of Rs 4 and a maximum upfront price of Rs 7.5. Similarly the final price ranges from Rs 4 to Rs 6.75. The variance for the plain pens is similar (see also graph 1).

In the second part of Panel A we now compare the demanded upfront payments by the wholesaler across printed and plain pens. The wholesalers, on average, initially ask for an upfront payment of 59% in the case of printed pens but ultimately accept an upfront payment rate of 36%. It is interesting to see that very few wholesalers demand a 100% upfront payment. In the case of plain pens, wholesalers on average start with an upfront payment demand of 19% and ultimately accept an average upfront payment of 12%. Wholesalers demand a substantially higher upfront payment for printed pens than for plain ones. These differences in the upfront payment rate between printed and plain pens are statistically significant and make sense intuitively since the wholesaler faces more risk of hold-up in case of printed pens. But again we see that there is wide variation around the mean (see also graph 2).¹⁵ An important point to note is that even within the same wholesaler we find wide variation in the upfront payment demanded across visits. Thus each wholesaler does not demand the same upfront payment across shoppers based on whether the order is printed or plain. In effect, there is significant variation in the contract structures offered to different shoppers visiting the same wholesaler.

¹⁴ Results are not reported but can be provided by the authors on request.

¹⁵ Table 2, Panel B reports the summary statistics for the renegotiation sample.

5. Ex-ante Negotiation and Relationship Specific Investments

In a first step we now want to establish whether there are differences in the contractual structure that are negotiated upfront between printed and plain pen purchases. If indeed printed pens require a bigger relationship specific investment as we argued above, wholesalers should expect a higher possibility of hold-up ex-post and thus undertake steps to alleviate these risks. The most extreme form of hold-up would be that the shopper does not return to pick up the order. Ideally wholesalers would then require an upfront payment of 100% to eliminate any risk of ex-post hold-up. However, this does not appear to be optimal since both parties face a two-sided hold-up problem. If the shopper was to pay the full amount upfront, the wholesaler could in turn engage in hold-up and not deliver the goods on time or indeed not deliver at all. Due to this mutual hold-up risk, both parties could be exposed to some ex-post renegotiation. However, it is safe to assume that the wholesaler has a higher reputation risk than the shopper.

To test how the specificity of the investment affects the upfront negotiation we look at two dimensions: (1) the fraction of upfront payment that is demanded between plain and printed pens and (2) the average price charged. In Table 3 we first test whether there are significant differences in the prices that are charged for printed versus plain pens. As discussed before we would expect there to be a mechanical difference due to the cost of printing. However, beyond the cost difference it will be interesting to see if the market incorporates any potential hold-up risk into the price, or whether wholesalers only use the upfront payment as a means to mitigate this risk. Columns (1) to (3) of Table 3 regress the initial price that is offered by the wholesaler on a dummy for whether the visit was assigned plain or printed pens and controls for other characteristics of the visit such as quantity purchased (log quantity), location of the store and pen brand. Location can be important if some wholesalers are in more highly frequented areas and thus have different demand patterns than those in less central areas. In column (1) we estimate a significant difference between the printed and the plain pens, however the size of the coefficient is 0.4, which means the magnitude of the differences is exactly the cost of printing. We then repeat this regression but include shopper and wholesaler fixed effects in columns (2) and (3). The results show that the magnitude of the coefficient on the printed dummy does not change when these added controls are included. However, the adjusted R-squared goes up significantly when including shopper and wholesaler fixed effects. In fact, the highest explanatory power in explaining price variation comes from wholesaler fixed effects (the adjusted R-squared increases from 0.16 to 0.74). Thus a large fraction of the variation is not explained by the specificity of the investment but is systematically related to the shopper and wholesalers.

In columns (4) to (6) we repeat these regressions for the final price that the wholesaler and the shopper agree upon after negotiation. We again see that the coefficient on the printed dummy is 0.4 which means the price difference between printed and plain pens stays equal to the cost of the printing. In fact, when we regress the percentage difference between the first and final offer price for the pens on the print dummy, we do not find a significant difference (column (7)).

Overall we find that wholesalers do not charge a price premium for printed pens beyond the cost of the printing, which suggests that any possible increased hold-up risk from printed pens is not included in the price. This is sensible since any price that is agreed upfront can be renegotiated ex-post. The results also suggest that wholesalers are more concerned about hold-up by shoppers rather than asset redeployability, i.e. a shopper might go bankrupt so that the wholesaler is stuck with the merchandise. In the latter case wholesalers should charge a higher risk premium on average for printed pens in order to insure themselves against the higher loss in case of default. However, if wholesalers believe that the more prevalent problem is one of strategic behavior - hold-up - then negotiating a higher price ex-ante does not help.

A more effective mechanism to limit the amount of ex-post hold-up in the wholesale market for pens is to increase the fraction of upfront payment that is requested from the shopper.¹⁶ Therefore in Table 4 we estimate the difference in the upfront payment that is demanded in the case of printed versus plain pens. The regressions in this table follow the setup in Table 3. In column (1) we use the fraction of upfront payment that is initially demanded by the wholesaler as the dependent variable and regress it on the print dummy and controls for the size of the order, the location and pen brand. The estimated coefficient on the print dummy is 0.39 and strongly significant. Also as compared to the price regressions reported in Table 3, the print dummy has a very high explanatory power for the upfront payments (adjusted R-squared of 0.31). This result is virtually unchanged if we include shopper and wholesaler fixed effects in columns (2) and (3). We then repeat these regressions for the final fraction of upfront payment that is agreed upon after bargaining between the shopper and the wholesaler, see columns (4) to (6). The coefficient on the print dummy now is a little smaller, 0.243, but still economically and statistically highly significant. When we regress the difference in the upfront versus final payment on the printed dummy in column (7) we find that there is a marginally significant negative coefficient. This suggests that wholesalers are more likely to be bargained down on the upfront payment in the case of plain pens, but are less willing to be bargained down to a smaller upfront payment in the case of printed pens.

From Table 3 and 4, we find that wholesalers do not charge a higher price (beyond the cost of printing) for printed pens to protect themselves against the risk of hold-up. The main mechanism used to mitigate hold-up risk is to charge a higher upfront payment. However interestingly, we find that the average magnitude of upfront (36% for printed pens) covers less than 50% of the wholesalers' cost price.¹⁷ This suggests that the wholesalers do not expect a high prevalence of hold-up in the market else one would expect the upfront payment to be much higher.¹⁸ In addition, from Table 4, we find that shopper fixed effects are much more important in explaining upfront payment as

¹⁶ One could argue that the upfront for printed pens is higher to cover the printing costs of the wholesaler, however, we find similar results even if we adjust the upfront for printing charges (if we remove printing charges from the final upfront, we get a difference of 0.17 on the print dummy as against 0.24 with it).

¹⁷ The profit margin of wholesalers is on average 15%.

¹⁸ One could argue that wholesalers also engage in hold-up and ex-post charge a higher price from shoppers in some cases and thus break even (see Hart and Moore, 1988). However, this seems highly unlikely as in all the transactions we conducted, we did not find a single instance where we were asked by the wholesaler to pay a higher price ex-post.

compared to the price dimension. We also find that there is wide variation in the level of upfront demanded by wholesalers across buyers. Thus, wholesalers must operate under the belief that there is a significant fraction of buyers in the market who do not engage in hold-up. In effect, competition forces wholesalers to understand the entire type distribution of buyers in order to set the upfront payment. If the fraction of buyers in the market who will engage in hold-up is low, then the level of upfront can be lower given the same price margin. In addition, wholesalers have to learn how to differentiate the types of shoppers who pose a threat of hold-up versus those that genuinely would never engage in hold-up. If a wholesaler develops a better screening mechanism, he cannot only demand a lower upfront, but could even lower the price margin to be more competitive.

Finally to understand how the requested fraction of upfront payment and the agreed-upon price relate to each other. In our bargaining script (and in the usual bargaining behavior of participants in the market) the parties first establish a price and then set the fraction of upfront payment. Therefore, we investigate whether the demanded upfront payment is a function of the price that was agreed upon.¹⁹ In Table 5 we regress the upfront payment on the price level in the overall sample and add an interaction between printed and plain pens. Columns (1) to (3) use the initial upfront payment as the dependent variable. We find a strong positive relationship between the price level and the fraction of upfront payment. The coefficient on the price is 0.182 and significant at the 1 % level. The magnitude and significance of the coefficient does not change even if we include shopper and wholesaler fixed effects in columns (2) and (3). In columns (4) to (6) we then repeat the same regressions for the final percentage of upfront payment. This might indicate that wholesalers have to be more concerned that a shopper reneges in a case where the bargaining led to a very high price.

Yet these results could also be interpreted as stemming from shopper and wholesaler specific fixed effects. While we are controlling for shopper and wholesaler fixed effects in these regressions, it is still possible that in some interactions between the two parties the shopper was more aggressive and thus got a better deal on all margins and vice versa for the wholesaler. But when we repeat the regressions for printed versus plain pens separately we see that the positive relation between upfront payment and price level is only significant for the plain pens. This result is interesting since it suggests that wholesalers attempt to “lock in” a plain pen deal with a high upfront when they are taking advantage of the shopper. The idea is that the shopper might realize ex-post that he got a bad deal. Thus without the high upfront payment they would not come back to complete the deal. This effect is not significant in the case of printed pens, which suggests that for printed pens the wholesalers are generally more careful in setting the upfront payment to minimize ex-post hold-up (independent of the price).

6. Ex-post Re-Negotiation

We now turn to analyze the dynamics of ex-post contract renegotiation. For that purpose we did a final visit to 75 randomly selected wholesalers. We conducted these additional

¹⁹ Note that in the re-negotiation sample (reported in Table 2, panel b), where we randomized the upfront payment to be paid by the shopper, we still find that price of printed pens is Rs 0.4 higher than plain pens, suggesting that wholesalers first set the price and then negotiate the upfront payment.

visits separately from the ex-ante experiment to avoid endogeneity in the upfront bargaining and the ex-post renegotiation. In particular, to understand the relationship between the renegotiation outcome and the upfront payment, we gave each shopper an effective limit for the maximum price and upfront payment they had to negotiate ex-ante. We did not let shoppers bargain down the wholesalers endogenously, since shopper-wholesaler pairs who have lower upfront payments up front may also have a different renegotiation interaction.

After an initial visit where the shopper placed an order and agreed on a contract and delivery date, the shoppers called up the wholesaler on the day of delivery to let him know that he cannot honor the order in full. The explanation for the renegotiation was that the shopper's ultimate customer had just canceled the order and the shopper would not be able to take possession of the order. The shopper would then offer to buy the order at a reduced price. We chose this "cover story" rather than asking the shoppers to hold-up the wholesaler outright, since we learned from focus group interviews with market participants that no one ever directly admits to this tactic but finds a "cover story" that sounds less aggressive. We wanted to avoid bargaining behavior that is out of the ordinary or could induce outright hostility by the wholesaler. For the same reason we also chose to conduct this renegotiation over the phone.²⁰ However, it is clear that the wholesalers understand that they are being held up since the wholesalers are asking to speak to the client or for the shopper to provide evidence that the client is canceling the order. After the renegotiation by phone had been finished, the shopper would call the wholesaler back after an hour and announce that he will pay the original contract in full since the final customer had just announced that they would honor the order after all.²¹ We decided to conduct the renegotiation only in the last visit to a wholesaler in order not to risk suspicion by the wholesaler and jeopardize the experiment.

Also, all the renegotiations were carried out over a 2-day period to avoid any risk of information spillover in the market.

6.1. Descriptive Statistics

Panel A of Table 6 shows that in 45% of the cases the price of the order is renegotiated. In the remaining 55% the wholesaler either refused to change the price or did not even agree to start a conversation about the price. In line with the predictions of the hold-up theories, we find that the willingness to renegotiate is much higher for printed as compared to plain pens (0.6 versus 0.32). While one would expect that wholesalers should always be willing to give a discount, given that the salvage value of printed pens is close to zero, in 14 cases of printed pens we find that wholesalers are unwilling to provide any reduction at all. Furthermore, in 7 cases that did not get renegotiated, the wholesaler hangs up on the shopper and implicitly adopted a take-it-or-leave-it attitude. These results suggest that some wholesalers might be concerned with reputation

²⁰ Note that our study did not lead to any distortions in the market equilibrium. Wholesalers did not experience a surge in the demand or price, on average. In addition, we also did not find any change in the upfront demanded by wholesalers after the study.

²¹ It is interesting to note that even in the cases where our shoppers tried to engage in hold-up but ultimately indicated that they need the pens, none of the wholesalers tried to hold up our shoppers when they came to pick up the pens and pay the remainder of the contract.

spillovers in the market or might feel aggrieved due to renegotiation of the contract by the shopper. Overall these incidences suggest that there are cases where renegotiation does not lead to the ex-post efficient outcome. In Panel B we now report whether the wholesalers agreed to give back the upfront payments when the renegotiation failed. We see that in none of the cases with printed pens did the wholesaler offer to return the upfront payment, but in 15 out of the 40 plain pen visits the wholesaler offered to return the upfront.²² In fact in 12 cases the wholesaler offered to give back the upfront in cash and in 3 cases they offered to give it back in kind, i.e. allow the shopper to buy something else for the equivalent amount of the upfront payment.

Finally, the first two rows of Panel C show the price difference that is achieved across all renegotiation visits, including those where the renegotiation failed. In the case of a failed renegotiation, we code the price change as zero. This is a strong assumption since it presumes that in case of failed renegotiation the shopper will pay the full price for the goods, which of course is not guaranteed. In fact this assumption clearly understates the size of the loss that a wholesaler would have to bear in a real world context since the likelihood of the shopper returning to pick up the goods is much lower in the case of the printed goods. But in that case the wholesalers' resale value of the pens is close to zero, while in the case of the plain pens the resale value is unaffected. In the next two rows of Panel C we describe the difference in price between the printed versus plain pens. Printed and plain pens on average have a mean price after renegotiation of Rs. 4.5 each. However, this masks a serious difference since the cost of printing in the printed pens is about Rs 0.4 per pen; so in fact the printed pens get renegotiated to a much lower price than the plain ones.

In line with the intuition of standard hold-up models, the results suggest that the bargaining position of a wholesaler is much lower in the case of printed pen orders since the outside value of the pens is zero. In fact the wholesalers themselves mention in the renegotiation that they cannot use the printed pens for anything now that they have been customized. This result is also supported by the fact that the minimum renegotiated price for the plain pen does not fall below the minimum price at the contracting stage, however it does for printed pens. Thus wholesalers are prepared to make a real loss on the pens, since they understand that the outside option is nearly zero.²³

6.2. Renegotiation Outcomes

To test the results of renegotiation more rigorously in column (1) of Table 7, we regress the renegotiation percentage on the printed pen dummy and shopper fixed effects. We cannot include wholesaler fixed effects in these regressions since we only have one visit to each of the wholesalers. The coefficient on the printed dummy is economically large (0.03) and statistically significant at the 1% level. In the renegotiation of printed pens, shoppers get a three percentage point higher discount on average relative to plain pens.

²² Note that we do not consider two cases of failed renegotiation of plain pens in this sample as the upfront paid in these cases is zero.

²³ We also find that the time spent in renegotiation is higher in case of printed pens. Refer to appendix A for the table.

This result confirms that wholesalers are well aware of the outside option they can charge for the goods, and once the pens have been printed, their outside value is very low.²⁴

In column (2) of Table 7 we also include observable characteristics of the deal such as the upfront payment that was made at the time of contract negotiation, and the price at which the deal was contracted. The coefficient on Final Upfront % is negative and significant, suggesting that wholesalers are less willing to renegotiate when they might feel that they have more bargaining power, since the shopper already has paid down a large fraction of the money. In contrast the coefficient on Price contracted (the price that was negotiated upfront) is positive, which suggests that wholesalers are more likely to give a discount if the original price left them with a higher profit margin. We can infer from this result that the wholesaler's bargaining power is higher in situations where a higher upfront payment was charged.

In columns (3) and (4) of Table 7 we now break out the sample into the renegotiations for printed versus plain visits. It is interesting to see that the coefficient on the upfront payment is negative, and significant in the case of printed pens (-0.195), but is not significant and close to zero in the case of the plain pens (-0.054). If we believe that the amount of upfront payment affects the shopper's ex-post bargaining power, we should expect that there is less renegotiation in cases where the shopper has a higher upfront payment, since the shopper now has less "room" to hold-up the wholesaler. It is interesting to see that this logic holds exactly in the case of printed pens, however, there is no effect in the case of plain pens. This suggests that the wholesalers perceive their bargaining power unaffected by the shoppers' demand in the case of plain pens, since their outside option is very high.

A second dimension that might affect the renegotiation is the price that was contracted upfront. In columns (3) and (4) of Table 7, we find a significant and positive relation between the percentage discount that the wholesaler agrees to in the renegotiation and the price that is contracted upfront. However, this positive relation only holds for the sample of plain pens but is not significant in the sample of printed pens. This implies that wholesalers are only willing to renegotiate the price ex-post if they were able to charge a high price for the plain pens upfront. So they might be willing to reduce the profit margin only if they are starting from a high margin. The same is not true for printed pens which implies that the wholesalers' willingness to reduce the price is independent of the initial contracted price. It suggests that the wholesaler understand that his bargaining power is very low in the case of printed pens and hence is willing to renegotiate the price across all printed pen deals and not only those that start with a high profit margin.

6.3. Reverse Hold-up

The last two columns of Table 7 focus on the likelihood that the shopper will get the upfront payment back during the renegotiation. As we saw before in the descriptive statistics in Table 6, this is only offered in case of plain pens. It is an interesting outcome

²⁴ In the renegotiation, the shoppers did not push the wholesaler very hard on the magnitude of the discount. Thus the renegotiation percentage is more likely to be a lower bound.

in itself that in none of the printed pen orders do the wholesalers offer to return the upfront, but they do in 15 out of 40 plain pen orders. Therefore we only focus on the set of plain pen renegotiations. This is a very interesting result since it suggests that a large fraction of wholesalers choose not to hold-up the shopper (by withholding the upfront payment) when they have the chance (and possibly the moral high ground) to do so.

In Columns (5) and (6) of Table 7 we regress a dummy for whether the wholesaler offered to refund the upfront payment on the deal characteristics, such as Final upfront %, Price contracted, Quantity of pens and Pen brand. In the regression in Column (5) we include only cases where the wholesaler did not offer to change the price. In Column (6) we repeat this regression for the sample of all plain deals. The second sample includes those deals where the wholesaler was willing to reduce the upfront price and code these as a zero, i.e. the shopper did not get the upfront payment back. The idea is that the wholesaler instead accepts a lower margin than letting the deal go altogether. This second coding is a little more ambiguous since one could also argue that it could have been coded as a one since the wholesaler does not take the upfront of the shopper and refuse to deliver the pens. We find a strong negative correlation with the amount of upfront payment and a somewhat weaker but negative coefficient with the price that was contracted upfront. In column (5) the coefficient on the upfront payment is -4.23 and is significant at the 1% level. The coefficient on negotiated price is -0.56 but only marginally significant.

These results suggest that the wholesaler is less willing to hold the shopper up even if he has the power to do so. However, it seems that the wholesaler does take into account the costs of being magnanimous: We find that wholesalers are not willing to return the upfront in the case of printed pens where they stand to make a bigger loss. Moreover, they are also less likely to return the upfront when the upfront payment itself was smaller. These results are consistent with wholesalers who trade off the possibility of building a reputation with the shopper against the money they give up for it. Interestingly the wholesalers are also slightly less likely to return the upfront if the upfront contracted was higher. One could argue, if wholesalers respond to financial incentives in their willingness to return the upfront, the underlying cause for these actions is most likely a rational calculation of the benefits of investing in reputation. If the underlying cause was social norms we might expect that the behavior is unaffected by the economic impact of the actions.

7. Conclusion

This paper uses a novel audit study methodology where we send trained auditors to execute real purchase orders to understand how relationship-specific investments affect the ex-ante contract structure and ex-post resolution of the ensuing hold-up problem.

We find that wholesalers use upfront payment as a mechanism to alleviate the risk of ex-post hold-up in line with the predictions of incomplete contracting models. Therefore, the upfront payment acts as a way of allocating bargaining power between contracting parties. We also find that there is a lot of variation in the contract structure that is not

explained by the relationship specificity of the investment. When looking at the ex-post contract renegotiation, we see that the wholesalers are more likely to renegotiate and accept lower price for printed pen purchases as compared to plain pens. But wholesalers offer a lower discount if they received higher upfront payments, especially in the case of printed pens, which again suggests that the allocation of upfront payments affects the bargaining outcome. But very surprisingly wholesalers are often willing to refund the upfront paid in cash (in the case of plain pens) even when the shopper holds up the wholesaler ex-post, even though they would be better off keeping the money.

On the one hand, these results confirm that contracting parties are concerned about ex-post hold-up when writing contracts and also at the stage of renegotiation. But on the other hand they also raise several interesting issues. The fact that wholesalers often refund the upfront payment or that they charge low upfront payments indicates that market participants do not expect high incidences of ex-post hold-up even when the counterparty has the possibility. In fact we see that wholesalers themselves do not fully exploit their bargaining power even when they have the opportunity (or even the right) to do so.

Then the question that arises naturally is whether these results reflect the prevalence of social norms that restrict wholesalers and shoppers from engaging in opportunistic behavior or whether reputational concerns prevent market participants from acting blatantly against the agreed contract. While the current paper cannot resolve this question in a definitive way, we provide contextual evidence that both types of constraints might be at play. We saw that wholesalers take into account the financial costs of being generous, i.e. their willingness to return the upfront payment depending on its size. One could argue that the underlying cause for these actions is most likely a rational calculation of the benefits of investing in reputation. If the underlying cause was social norms we might expect that the behavior is unaffected by the economic impact of the actions. At the same time we also observe that wholesalers expect low hold-up probabilities even from those shoppers who they have never interacted with before and they do not know from any social or professional context, so the reputational implications are minimal. So the prevailing social norm seems to be one of honoring contracts even if the enforcement is incomplete. To shed more light on these questions we believe that a natural next step for further research is to explore how contextual and economic forces mitigate or exaggerate the risk of hold-up.

References

- Acharya, V., Bharath, S., Srinivasan, A., 2007. "Does Industry-wide Distress Affect Defaulted Firms? - Evidence from Creditor Recoveries," *Journal of Financial Economics*, 85, 787-821
- Aghion, P., and Bolton, P., 1992. "An Incomplete Contracts Approach to Financial Contracting," *Review of Economic Studies*, 59, 473-494.
- Ayres, I. and Siegelman, P., 1994 "Race and Gender Discrimination in Negotiation for the Purchase of a New Car," *American Economic Review*, 85, 304-321.
- Banarjee, A., Duflo, E., 2000. "Reputation Effects and the Limits of Contracting: a study of the Indian Software Industry," *Quarterly Journal of Economics*, 115, 989-1017.
- Benmelech, E., Bergman, N., 2008. "Liquidation Values and the Credibility of Financial Contract Renegotiation: Evidence from U.S. Airlines," *Quarterly Journal of Economics*, Forthcoming.
- Bertrand, M., Mullainathan, S., 2004. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review*, 991-101.
- Bolton, P., and Scharfstein, D., 1996. "Optimal Debt Structure with Multiple Creditors," *Journal of Political Economy*, 104, 1-26.
- Camerer, C., and Thaler, R., (1995). "Ultimatums, Dictators, and Manners," *Journal of Economic Perspectives*, 9, 209-219.
- Crocker, K. J., Reynolds, K. J., 1993. "The Efficiency of Incomplete Contracts: An Empirical Analysis of Air Force Engine Procurement," *Rand Journal of Economics*, 24, 126-146.
- Fehr, E., and Schmidt, K. M., (1999). "A Theory of Fairness, Competition and Co-Operation," *Quarterly Journal of Economics*, 114, 817-868.
- Grossman, S., Hart, O., 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 691-719.
- Güth, W., Schmittberger, R., and Tietz, R., (1990), "Ultimatum Bargaining Behavior—A Survey and Comparison of Experimental Results," *Journal of Economic Psychology*, 11, 417-449.
- Hart, O., Moore, J., 1988. "Incomplete contracts and Renegotiation," *Econometrica of Political Economy*, 56, 755-785.

Hart, O., Moore, J., 1990. "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98, 1119-1158.

Joskow, P., 1985. "Vertical Integration and Long-Term Contracts: The Case of Coal-Burning Electric Generation Plants." *Journal of Law, Economics, and Organization*, 1, 33-80.

Joskow, P., 1987. "Contract Duration and Relationship-Specific Investment: Empirical Evidence from Coal Markets." *American Economic Review*, 77, 168-85.

Kaplan, S. N., Stromberg, P., 2003. "Financial Contracting Theory Meets the Real World: An Empirical Analysis of Venture Capital Contracts," *Review of Economic Studies*, 70, 281-315.

Klein, B., Crawford, R.G., and Alchian, A. A., 1978. "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process," *Journal of Law & Economics*, 21, 297-326.

McMillan, J., Woodruff, C., 1999. "Interfirm Relationships and Informal Credit in Vietnam," *Quarterly Journal of Economics*, 114, 1285-1320.

Newmark, D., Bank, R.J., and Van Nort, K.D., 1996. "Sex Discrimination in Restaurant Hiring: An Audit Study", *The Quarterly Journal of Economics*, 111, 915-941.

Pulvino, T. C., 1998. "Effects of Bankruptcy Court Protection on Asset Sales," *Journal of Financial Economics*, 52, 151-86.

Roth, A. E., 1995. "Bargaining Experiments," in J. Kagel and A. Roth, eds., *Handbook of Experimental Economics* (Princeton: Princeton University Press).

Shleifer, A., Vishny, R.W., 1992. "Liquidation Values and Debt Capacity: A Market Equilibrium Approach," *Journal of Finance*, 47, 143-66.

Williamson, O., 1979, "Transaction Cost Economics: The Governance of Contractual Relations," *Journal of Law and Economics*, October 1979, 22, 233-261.

Williamson, O., 1985, *The Economic Institutions of Capitalism*, New York, NY: Free Press.

Williamson, O., 1988, "Corporate Finance and Corporate Governance," *Journal of Finance*, 43, 567-592.

Table 1: This table presents the summary statistics of the visits. Panel A reports the summary statistics of no of visits to each wholesaler and by each shopper. Average order size per visit is the average no of pens purchased per visit. Panel B presents the summary statistics for visits of where the contract was renegotiated ex-post.

Panel A	Obs	Avg no of visits	Median no of visits	Std. Dev	Min no of visits	Max no of visits	Average order size per visit
no of wholesalers	107	4.61	5	1.05	2	6	615
no of shoppers	46	10.73	11	0.90	6	13	617

Panel B (Renegotiation sample)	Obs	Avg no of visits	Median no of visits	Std. Dev	Min no of visits	Max no of visits	Average order size per visit
no of wholesalers	75	1	1	0	1	1	689
no of shoppers	15	5.4	5	1.585	3	9	686

Table 2: This table presents the summary statistics of the price and the upfront payment % demanded during the visits. Initial offer is the initial price per pen (Rs.) offered by the trader. Printed pen refers to pen on which a shopper gets a customized message printed. Final rate is the final contracted rate per pen (including printing costs if any). Initial upfront % is the initial advance payment demanded by the wholesaler as a fraction of total sales price. Final upfront % is the final advance paid as a fraction of total sales price. Panel B presents the summary statistics for visits where the contract was renegotiated ex-post.

Panel A


Price	Obs	Mean	Median	Std. Dev	Min	Max
Initial offer-Printed pen	240	5.33	5.22	0.772	4.0	7.5
Final rate-Printed pen	240	4.93	4.82	0.582	4.0	6.75
Initial offer-Plain pen	254	4.90	4.8	0.665	3.8	6.5
Final rate- Plain pen	254	4.53	4.45	0.465	3.8	6
Upfront payment						
Initial upfront %-Printed pen	240	0.588	0.5	0.299	0	1
Final upfront %-Printed pen	240	0.360	0.32	0.223	0	1
Initial upfront %-Plain pen	254	0.192	0	0.286	0	1
Final upfront %-Plain pen	254	0.117	0	0.202	0	1

Panel B (Renegotiation sample)


Price	Obs	Mean	Median	Std. Dev	Min	Max
Initial offer-Printed pen	35	5.15	5.25	0.606	4.0	6.5
Final rate-Printed pen	35	4.83	4.9	0.438	4.0	5.9
Initial offer-Plain pen	40	4.75	4.65	0.627	3.85	6.0
Final rate- Plain pen	40	4.42	4.30	0.410	3.8	5.25
Upfront payment						
Initial upfront %-Printed pen	35	0.526	0.5	0.269	0	1
Final upfront %-Printed pen	35	0.232	0.24	0.087	0.08	0.41
Initial upfront %-Plain pen	40	0.230	0.15	0.284	0	1
Final upfront %-Plain pen	40	0.091	0.06	0.068	0	0.27

Graph 1

Final rate per pen (Printed sample)




Final rate per pen (Plain sample)



Graph 2

Final Upfront % (Printed sample)



Final Upfront % (Plain sample)




Table 3: Regressions of Initial price offered and Final price contracted

This table reports the results of OLS regressions. The dependent variables are Initial price offered per pen, Final contracted rate per pen (including printing costs if any) and Price diff where Price diff is defined as (Initial price offered –final price contracted)/ Initial price offered. Print is a dummy variable that takes the value of one if customized printing was done on the pen. Quantity is the log of the size of order. Brand is a dummy for the type of pen (we have 2 brands of pen that we purchase). Location is a dummy variable that takes the value of one for wholesalers that are not located in the main street. Shopper fixed effects refer to fixed effects for each individual shopper. Wholesaler fixed effects refer to fixed effects for each individual wholesaler. White heteroskedasticity consistent standard errors are reported in parentheses. The symbols ***, **, * indicate significance levels of 1%, 5%, and 10% respectively.

	Initial price offered			Final price contracted			Price diff
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Print	0.415*** (0.064)	0.414*** (0.064)	0.396*** (0.037)	0.399*** (0.047)	0.405*** (0.048)	0.393*** (0.027)	-0.004 (0.004)
Quantity				-0.255 (0.162)			
Pen brand				-0.143** (0.064)			
Location				0.204* (0.107)			
Constant	5.300*** (1.473)	4.911*** (0.041)	4.675*** (0.190)	6.010*** (1.055)	4.531*** (0.029)	4.234*** (0.158)	0.086*** (0.022)
Shopper fixed effect	no	yes	yes	no	yes	yes	yes
Wholesaler fixed effect	no	no	yes	no	no	yes	yes
N	494	494	494	494	494	494	494
Adj-R2	0.095	0.160	0.741	0.144	0.178	0.776	0.300

Table 4: Regressions of Initial upfront offered and Final upfront contracted

This table reports the results of OLS regressions. The dependent variables are Initial upfront % offered, Final upfront % contracted and Upfront diff where Upfront diff is defined as (Initial upfront % offered – final upfront % contracted)/ Initial upfront % offered. Print is a dummy variable that takes the value of one if customized printing was done on the pen. Quantity is the log of the size of order. Brand is a dummy for the type of pen (we have 2 brands of pen that we purchase). Location is a dummy variable that takes the value of one for wholesalers that are not located in the main street. Shopper fixed effects refer to fixed effects for each individual shopper. Wholesaler fixed effects refer to fixed effects for each individual wholesaler. White heteroskedasticity consistent standard errors are reported in parentheses. The symbols ***, **, * indicate significance levels of 1%, 5%, and 10% respectively.

	Initial upfront % offered			Final upfront % contracted			upfront diff
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Print	0.396*** (0.026)	0.385*** (0.026)	0.375*** (0.026)	0.246*** (0.019)	0.244*** (0.017)	0.231*** (0.017)	-0.059* (0.033)
Quantity				-0.037 (0.063)			
Pen brand				0.044** (0.019)			
Location				-0.018 (0.031)			
Constant	0.497 (0.599)	0.197*** (0.017)	0.385** (0.158)	0.357 (0.408)	0.116*** (0.012)	0.394*** (0.147)	0.220** (0.110)
Shopper fixed effect	no	yes	yes	no	yes	yes	yes
Wholesaler fixed effect	no	no	yes	no	no	yes	yes
N	494	494	494	494	494	494	343
Adj-R2	0.314	0.418	0.511	0.257	0.392	0.526	0.190

Table 5: Regressions of Initial upfront % offered and Final upfront % contracted on Final price contracted

This table reports the results of OLS regressions. The dependent variables are Initial upfront % offered, Final upfront % contracted. Print is a dummy variable that takes the value of one if customized printing was done on the pen. Price contracted is the final contracted price per pen (including printing costs if any). Shopper fixed effects refer to fixed effects for each individual shopper. Wholesaler fixed effects refer to fixed effects for each individual wholesaler. White heteroskedasticity consistent standard errors are reported in parentheses. The symbols ***, **, * indicate significance levels of 1%, 5%, and 10% respectively.

	Initial upfront % offered			Final upfront % contracted			Final upfront % contracted	
	(1)	(2)	(3)	(4)	(5)	(6)	Printed sub-sample	Plain sub-sample
Print	0.813*** (0.247)	0.931*** (0.235)	1.216*** (0.236)	0.338* (0.199)	0.449*** (0.174)	0.671*** (0.167)		
Price contracted	0.172*** (0.041)	0.183*** (0.038)	0.253*** (0.065)	0.114*** (0.032)	0.115*** (0.028)	0.173*** (0.048)	0.036 (0.074)	0.213* (0.108)
Print*Price contracted	-0.098* (0.053)	-0.125*** (0.050)	-0.190*** (0.050)	-0.028 (0.043)	-0.051 (0.037)	-0.103*** (0.036)		
Constant	-0.588*** (0.184)	-0.635*** (0.172)	-0.700** (0.323)	-0.402*** (0.143)	-0.406*** (0.128)	-0.349 (0.262)	0.467 (0.373)	-0.401 (0.488)
Shopper fixed effect	no	yes	yes	no	yes	yes	yes	yes
Wholesaler fixed effect	no	no	yes	no	no	yes	yes	yes
N	494	494	494	494	494	494	240	254
Adj-R2	0.347	0.441	0.533	0.292	0.423	0.545	0.472	0.274

Table 6: This table presents the summary statistics of the visits where the contract was renegotiated ex-post. Panel A, decomposes the number of visits where the renegotiation succeeded (price was reduced) versus visits where renegotiation failed (the price was not changed) by printed/non-printed. Renegotiation Abruptly terminated refers to visits where the wholesaler abruptly ended the negotiation with the shopper (wholesaler abruptly hung-up the phone). Panel B decomposes the number of visits where renegotiation failed (the price was not changed) by whether the wholesaler agreed/refused to refund the upfront paid. Also, it reports in case of refund, whether the refund was in cash or in kind (in form of purchase of other items). Panel C reports the summary statistics of Renegotiation percentage by type of visit (printed/non-printed). Renegotiation percentage is defined as the (final contracted price per pen-price per pen after renegotiation)/ final contracted price per pen (note that final contracted price per pen is the price per pen that was agreed before renegotiation). Panel C also presents the price per pen after renegotiation for the sub-sample where the renegotiation succeeded. Printed pen refers to pen on which a shopper gets a customized message printed.

Panel A

	Obs	Renegotiation Succeeded	Renegotiation Failed	Renegotiation Abruptly terminated	Fraction Renegotiated
Printed pens	35	21	14	8	0.6
Plain pens	40	13	27	7	0.32
Total visits	75	34	41	15	0.45

Panel B

	No of Failed renegotiation visits	Upfront refunded (in case of renegotiation failure)	Refund in cash	Refund in kind
Printed pens	14	0	0	0
Plain pens	25	15	12	3

Panel C

Renegotiation percentage	Obs	Mean	Median	Std. Dev	Min	Max
Printed pens (including printing costs)	35	0.041	0.030	0.048	0.00	0.216
Plain pens	40	0.012	0.000	0.020	0.00	0.061
Price after renegotiation	Obs	Mean	Median	Std. Dev	Min	Max
Printed pens (including printing costs)	21	4.49	4.5	0.383	3.8	5.2
Plain pens	13	4.50	4.5	0.352	4.0	5.1

Table 7: Regressions of the magnitude of renegotiation percentage and Likelihood of getting the Upfront paid refunded

This table reports the results of OLS regressions in column 1-4. The dependent variable is Renegotiation percentage where renegotiation percentage is defined as the (final contracted price per pen-price per pen after renegotiation)/ final contracted price per pen (note that final contracted price per pen is the price per pen that was agreed before renegotiation). Column 5-6 report results of a probit (the co-efficient reported are marginal effects). The dependent variable is Upfront refund dummy which takes the value of one if the wholesaler agrees to refund the Upfront paid and zero otherwise Column 5 reports the results for the sample where the price renegotiation failed. Column 6 reports the results for the entire sample. Print is a dummy variable that takes the value of one if customized printing was done on the pen. Final upfront % is the final advance paid as a fraction of total sales price. Price contracted is the final contracted price per pen (including printing costs if any). Quantity is the log of the size of order. Brand is a dummy for the type of pen (we have 2 brands of pen that we purchase). Location is a dummy variable that takes the value of one for wholesalers that are not located in the main street. Shopper fixed effects refer to fixed effects for each individual shopper. White heteroskedasticity consistent standard errors are reported in parentheses. The symbols ***, **, * indicate significance levels of 1%, 5%, and 10% respectively.

	Renegotiation percentage				Upfront Refund Dummy	
	(1)	(2)	Printed sample (3)	Plain sample (4)	Failed Renegotiation sample (5)	(6)
Print	0.028*** (0.008)	0.029** (0.013)				
Final Upfront %		-0.136** (0.065)	-0.195** (0.079)	-0.052 (0.039)	-4.229*** (1.894)	-3.275** (1.578)
Price contracted		0.021** (0.010)	0.022 (0.018)	0.028*** (0.007)	-0.564* (0.329)	-0.808*** (0.243)
Quantity		-0.032 (0.132)	0.139 (0.154)	0.094* (0.054)	0.709 (2.560)	-0.657 (1.615)
Pen brand		-0.040 (0.063)	-0.001 (0.061)	0.035 (0.022)	0.225 (1.026)	-0.290 (0.590)
Location		-0.002 (0.011)			-0.389 (0.320)	
Constant	0.012*** (0.003)	0.183 (0.863)	-0.935 (1.081)	-0.734** (0.359)		
Shopper Fixed effect	no	yes	no	no	no	no
N	75	75	35	40	25	37
Adj/Pseudo-R2	0.135	0.207	0.230	0.341	0.330	0.341

Appendix A: Regressions of the amount of time spent in renegotiation

This table reports the results of OLS regressions. The dependent variable is Renegotiation Duration where renegotiation duration is defined as the number of minutes spent in renegotiation. Print is a dummy variable that takes the value of one if customized printing was done on the pen. Shopper fixed effects refer to fixed effects for each individual shopper. White heteroskedasticity consistent standard errors are reported in parentheses. The symbols ***, **, * indicate significance levels of 1%, 5%, and 10% respectively.

	Renegotiation Duration	
	(1)	(2)
Print	1.628*** (0.571)	1.718*** (0.594)
Constant	3.200*** (0.275)	3.157*** (0.341)
Shopper Fixed effect	No	Yes
N	75	75
Adj-R2	0.106	0.137