THE LIMITS OF NEIGHBORLY EXCHANGE*

RYAN BUBB†, SUPREET KAUR‡, AND SENDHIL MULLAINATHAN∓

February 16, 2018

Abstract. Informal contracting among individuals underpins economic activity in developing countries. We design a simple test to detect failures in intertemporal trade among neighboring farmers in Indian villages. We offer to subsidize the cost of irrigation among buyer and seller pairs, and vary the seller's expected ability to ensure future receipt of funds: the subsidy payment is delivered into the hands of either the seller (Seller-subsidy) or buyer (Buyer-subsidy). Relative to the Seller-subsidy, the Buyer-subsidy results in 58% less irrigation and a 0.34 standard deviation decrease in the buyer's crop yields. These effects are not eliminated through experience or social or caste linkages. The surplus left on the table under the Buyer-subsidy corresponds to 16.1% of annual household income. These findings suggest that within the context of our experiment, barriers to interpersonal contracting have large consequences for investment, output, and earnings.

* This paper was previously circulated under the title "Contracting Failures in the Village Economy." We thank Siwan Anderson, Emily Breza, Patrick Francois, Rocco Machiavello, Bentley MacLeod, Jeremey Magruder, Ameet Morjaria, Jann Spiess, and Chris Woodruff for their helpful comments. We are grateful to the Institute for Financial Management and Research (IFMR) and ICICI Bank for assistance with field research. Dominik Bulla, Nicole Brunda, Sanchit Kumar, Medha Aurora, Arnesh Chowdhury, and Evan Plous provided excellent research assistance. This project received financial support from Harvard University, the IFC, and IFMR. Any errors are our own.
†School of Law, New York University. Email: ryan.bubb@nyu.edu.
‡Department of Economics, UC Berkeley and NBER. Email: supreet@berkeley.edu. (Corresponding author).
∓Department of Economics, Harvard University and NBER. Email: mullain@fas.harvard.edu.

1. Introduction

The economic lives of the poor are characterized by informal contracting among individuals. In village economies, this is true of the markets for credit, savings, insurance, land, labor, and capital inputs. For example, to the extent that the poor access credit or insurance, it is almost exclusively through neighbors, relatives, and local moneylenders (e.g. Banerjee and Duflo 2007, Collins et al. 2010). Similarly, the bulk of agricultural input transactions—e.g., for labor, land, irrigation, bullock rental, and farm machinery rental—are comprised of bilateral sales between individual co-villagers. The functioning of the contracting environment among individuals is therefore essential for investment and output in this setting.

In this paper, we empirically examine this contracting environment. While contracting has many facets, we focus on one feature with relevance for productive investment: intertemporal contracting between buyers and sellers. Specifically, when providing a productive input, the seller must bear the cost of the input now, while the payoff to the investment is realized by the buyer in the future (e.g., at harvest). If buyers cannot compensate sellers up front—for example, due to liquidity constraints—then mutually beneficial trade can still occur if the seller is willing to wait until the buyer's benefit is realized before collecting payment. However, if there are barriers to future funds recovery, this may prevent mutually beneficial trade from occuring in the first place. Such barriers could arise from simple default risk, which is inherent in any intertemporal exchange, or from features that may be more specific to the backdrop of the village economy, such as perceived social costs of enforcing collection from one's neighbors.

We design a simple test for whether the cost of future funds recovery undermines mutually beneficial trade. The backdrop for our test is village irrigation markets in India, in which smallholder farmers purchase water from a neighboring well owner each year.[1] To enable a test for enforcement, we introduce a cash subsidy for buyer-seller pairs. Specifically, if the seller delivers irrigation to the buyer during the hot season, then the pair receives the subsidy.[2] However, this subsidy is paid out well after the irrigation takes place—a month after the hot season ends.

This temporal lag in subsidy delivery gives us a lever with which to manipulate enforcement: we randomize in whose hands the future subsidy payment will be delivered. Specifically, in the Buyer-subsidy treatment, we tell the buyer-seller pair that the money will be delivered into the hands of the water buyer. In the Seller-subsidy treatment, the pair is told the money will be delivered directly to the water seller. The buyer and seller

[1]Because it is costly to transport water over long distances, water buyers can effectively only purchase water from a neighboring farmer whose land (and therefore well) is in close physcial proximity to their own. A given water buyer typically has 1-5 potential sellers from whom he could purchase water. Each buyer purchases water from his neighbors multiple times every year.

[2]The subsidy amounts to 50% of the market price of a typical irrigation.

then take this information into account when deciding whether to trade during the hot season. Note that the timing of events, information available to the parties, the amount of liquidity, and the total surplus from trade are all exactly the same across both treatment conditions. The only difference is into whose hands the subsidy money will arrive. In the first treatment condition, the seller must trust the buyer to transfer funds to him, whereas the second condition ensures the money arrives directly to the seller.

Under the Coasian benchmark of perfect enforcement, the buyer and seller will ex ante agree on how to split the subsidy payment when it arrives, and there should be no difference in the level of trade across the two subsidy treatment conditions. In contrast, if the seller cannot trust the buyer to transfer (some subset of the) future subsidy payment to him, then he may be unwilling to bear the cost of providing irrigation up front. Thus, in the presence of enforcement constraints, the amount of irrigation will be higher under the Seller-subsidy relative to the Buyer-subsidy.[3] In addition, by examining the resultant impact on output (i.e. yields), we can assess whether the magnitude of enforcement failures is economically meaningful.

While this design enables us to construct a clean test for enforcement, its interpretation requires an important caveat. Because the subsidy increases the gains from trade, it induces irrigations that may not be efficient in its absence. We therefore cannot presume that trades that occur under the Seller-subsidy should occur in equilibrium. However, our goal is not to assess the underlying efficiency of the irrigation market. Rather, we use this setting—which involves trade among neighbors who buy and sell irrigation multiple times every year—as a convenient backdrop for our test.

To implement the experiment, we identify water buyer-seller pairs across 21 villages in central Uttar Pradesh, India.[4] Among the 407 pairs in our sample, 94% of buyers purchased irrigation from a neighbor in the previous year, and 63% of pair members have traded with each other in the past. These pairs are randomized into one of three treatment groups: the Buyer-subsidy, Seller-subsidy, and a pure Control group where no subsidy is offered. The subsidy offer applies to irrigations during the hot season (April-June), and subsidy payments delivered in July.

The effects on the level of ex ante trade point to the presence of enforcement constraints. The probability that a buyer-seller pair trades in a given week is 3.6 percentage points (31%)

[3]As we discuss in Section 2 below, this prediction cannot be explained by alternate mechanisms such as different returns to capital among buyers and sellers, or differential trust among buyers and sellers that we will deliver the subsidy. In short, if there are gains from trade in the case of the Seller-subsidy treatment, then these gains exist in the Buyer-subsidy treatment; outcomes on the amount of trade should look the same as long as the seller can recover the subsidy from the buyer when it is delivered.

[4]The sample was constructed as follows. We identified farmers who own a plot of land without a well on it (water buyers). For each of these farmers, we identified all the well owners around their plot of land who were physically close enough to sell them water; from this group, we randomly chose one potential water seller.

higher under the Seller-subsidy relative to the Buyer-subsidy.[5] Overall, Seller-subsidy pairs engage in 58% more hours of irrigation relative to those in the Buyer-subsidy.

This difference in irrigation has a substantial impact on output. On average, water buyers have crop yields that are 0.335 standard deviations higher when the subsidy is delivered to the seller rather than to themselves. This corresponds to an estimated 9% increase in crop revenue. These findings indicate that within the specific context of our experiment, the enforcement failure among the Buyer-subsidy group (relative to the benchmark of the Suyer-subsidy) is economically meaningful.

We next turn to examine the details of contract terms to obtain supporting evidence for how enforcement constraints affect the structure of trade. We find that pairs seldom enter into contracts involving an ex post transfer of the subsidy. Specifically, among those pairs that receive a positive subsidy payment, in 87% of cases, subsidy recipients and their trading partners agree ex ante that the subsidy recipient will not transfer funds to the opposite party after the subsidy is delivered. This is consistent with anticipated costs to enforcing ex post contracts, leading people to not enter into such contracts in the first place. While only suggestive evidence, this is is in line with our hypothesis that the Buyer subsidy will be less successful at inducing ex ante trade because sellers do not trust buyers to transfer funds to them once the subsidy is delivered.

If ex post transfers are costly to enforce, then parties can use up front payments at the time of irrigation to enable trade to occur. Consistent with this, in the Seller-subsidy, sellers offer a price discount to encourage trade—in anticipation of the fact that they will receive a subsidy payment from us in the future. Specifically, in the Seller subsidy group, sellers are 8 percentage points (97%) more likely to give their paired buyer a price discount than those in the Buyer subsidy group. They are also more likely to offer trade short term credit—allowing buyers to repay after the irrigation date.

This highlights an important role for liquidity. In the presence of enforcement constraints, the buyer's wealth at the time of irrigation will constrain his ability to make ex-ante transfers to enable trade. Consistent with this, we find that when buyers are poorer, they are not able to take advantage of the Buyer subsidy treatment: trade under the Buyer subsidy is not statistically different than trade under the Control group of no subsidy. It is only when buyers are wealthier that we see significant increases in trade under the Buyer subsidy relative to Control. In contrast, sellers' wealth has no positive impact on trade under the Buyer subsidy.

Finally, we examine potential correlates of informal enforcement mechanisms—including previous trading relationships between the buyer and seller, market power, or being within

[5]We measure the amount of irrigation via weekly surveys. Our field staff visited each buyer and seller every week over the course of the 3 month irrigation period (hot season), and verified reports of irrigation by checking soil moisture on the buyer's plot. We also conducted an endline survey to obtain information on yields and other endline outcomes.

the same caste network. Overall, we find little evidence that correlates of relational contracting enable pairs to fully overcome the enforcement problem. Even in cases where social linkages are high, the Buyer subsidy group generally trades less than the Seller subsidy group. Overall, these results suggest that the participants in our study have limited ability to overcome the enforcement problem, at least within the context of our study.

While these correlates do not seem to enable pairs to overcome the enforcement problem, we do see evidence that they matter for contracting more generally. In the subsidy groups overall, 62% of pairs never engage in any trade. This suggests that for these pairs, a 50% subsidy was not sufficient to make trade worthwhile. Correlates of social and market linkages—e.g. whether the buyer and seller are the same religion or if they have traded in the past—have strong predictive power for which pairs end up taking advantage of the subsidy offer (versus having the same level of trade as the control group). These heterogeneous effects could just reflect efficient contracting. Alternately, they could be indicative of a more general fundamental contracting failure—with transactions costs increasing so much with social distance that even a 50% subsidy is not enough to induce trade on the margin. This is consistent with Anderson (2011), who presents evidence that caste differences are correlated with lower levels of irrigation transactions in Indian villages—with substantial implications for the yields of low caste farmers. While our experiment is not well suited to examining the underlying cause for this pattern of results, it constitutes an interesting direction for additional research.

This paper contributes to the empirical literature on contracting failures in developing countries. A growing number of studies points to the relevance of enforcement issues (McMillan and Woodruff 1999; Machiavello and Morjaria 2014, 2016; Machiavello and Miquel-Florense 2015). These papers apply clever identification strategies to contract data to examine the role of relational contracting among formal firms. They find support for the idea that repeat relationships play an important, though imperfect, role in enabling transactions among firms. [6]

Our study design complements these analyses. Using the Seller-subsidy case as a benchmark for how much irrigation could occur under the subsidy, we can quantify the extent to which the Buyer-subsidy group falls short. The effects on irrigation and yields indicate that the enforcement failure is large in magnitude. Of course, this magnitude applies specifically to the context of the trading opportunity we create through the subsidy offer. However, the fact that such large failures arise in our population—neighboring farmers who routinely buy and sell irrigation, and who have traded with each other in the past—suggests that enforcement problems may not be limited to just the narrow case of our subsidy offer. For example, these findings may be relevant in understanding why we rarely observe long-term

[6] In addition, a small set of studies explores other, but related, contracting issues. Banerjee and Duflo (2000) present evidence that reputation is used to solve incomplete contracting problems in the Indian software industry. Iyer and Schoar (2008) use a field experiment to examine how concerns about hold-up affect the timing of payments and ex-post barganing.

trade credit in irrigation markets in equilibrium.[7] More generally, these findings suggest that enforcement problems have the potential to hamper mutually beneficial trade in village economies, which are characterized by the prevalence of informal interpersonal exchange.

2. Model

2.1. **Set-up.** We construct a simple stylized model in which a buyer and seller have a bilateral trading opportunity. The buyer has access to a production technology. In period 1, the buyer can invest one unit of input, with deterministic revenue r realized in period 2. Denote the seller's cost of supplying the input as c; this cost is borne by the seller at the time of delivery (i.e. in period 1). Let w denote the buyer's wealth (i.e. cash on hand) at the beginning of period 1.

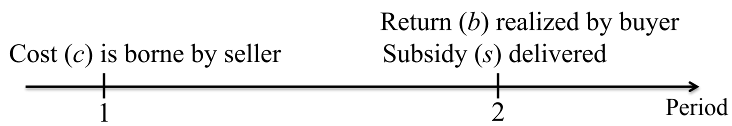

Figure 1. Model - Timing of Events

We introduce a subsidy payment, $s > 0$, that is delivered to one of the parties in period 2 *if* trade occurs in period 1. Without loss of generality, we ignore time discounting between periods for simplicity. With the subsidy, the total surplus from trade is therefore $r + s - c$.

The parties can write a contract in period 1 over payments in periods 1 and 2. However, a contract that requires the recovery of funds in period 2 incurs a potential loss, as detailed below. Throughout the model, we assume the seller recovers as much of the payment as possible from w at the time of sale in period 1, and recovers the rest in period 2.[8]

In what follows, we compare two cases. In the "Buyer Subsidy" case, the subsidy payment s is delivered into the hands of the *buyer* in period 2 if trade occurs in period 1. In the "Seller Subsidy" case, the subsidy payment s is delivered into the hands of the *seller* in period 2 if trade occurs in period 1. Our primary interest is in comparing how these cases impact whether trade occurs in period 1.

2.2. **Conditions for Trade.** For ease of exposition, let s_B denote the subsidy given to the buyer and let s_S denote the subsidy given to the seller, where $s_B + s_S \equiv s$. For example, in the Buyer Subsidy case, $s_B = s$ and $s_S = 0$. The opposite is true in the Seller subsidy case.

[7] For example, among the Control group (where no subsidy was offered), 97% of irrigation payments are due within one week of the irrigation date, and no sellers allow buyers to wait until harvest to repay.

[8] This assumption is innocuous. It simplifies exposition, without qualitatively altering the results. In the absence of intertemporal contracting frictions, this assumption makes no difference to the results: the parties will be indifferent to when funds are paid. In the presence of such frictions, this rule maximizes the seller's willingness to trade, since funds arrive up front at the time of sale.

Let p_1 and p_2 denote the payments made by the buyer to the seller in periods 1 and 2, respectively. The total payment that the seller receives for the input is therefore $p \equiv p_1 + p_2$. In each period, the payment made by the buyer must be weakly less than the buyer's cash on hand. In period 1, the seller can receive a payment amount up to w from the buyer:

$$p_1 \leq w.$$

In period 2, the buyer will have his remaining liquid wealth, $w - p_1$, plus his revenue r and the subsidy payment s_B. From the perspective of the seller in period 1, the maximum that can be recovered from the buyer in period 2 is:

$$E_1 [p_2] \leq \mu [r + (w - p_1)] + \lambda s_B,$$

where $\mu \leq 1, \lambda \leq 1$ denote the probabilities that the seller will be able to recover payment of non-subsidy income and subsidy income, respectively, in period 2. We allow each of these probabilities to differ, though it is possible that $\mu = \lambda$.

Recall our assumption above that the seller recovers as much of the payment as possible at the time of sale in period 1. Substituting for $p_1 = w$ in the second inequality and adding the two conditions together, the constraint on the maximum total payment that the seller can expect to recover reduces to:

$$E_1 [p] \leq w + \mu r + \lambda s_B. \tag{2.1}$$

This condition is intuitive. The seller can recover up to w with certainty (since the buyer has this much cash on hand at the time of sale in period 1). However, recovery of funds in period 2 incurs a potential expected loss. In expression (2.1), μ and λ reflect intertemporal contracting costs. We have interpreted μ and λ as the expected probabilities that the seller can enforce payment by the buyer in period 2. More generally, these parameters could reflect any type of cost that prevents period two recovery—for example, a personal or social cost of visiting a neighbor at home to collect money.[9]

Proposition 1: Conditions for Trade *For each subsidy case, trade will occur if and only if two conditions are satisfied:*

Buyer subsidy	Seller subsidy
(1) $b + s - c \geq 0$	(1) $b + s - c \geq 0$
(2) $w + \mu r + \lambda s \geq c$	(2) $w + \mu r + s \geq c$

The first condition is an efficiency condition: for there to be trade, the total surplus from trading must be nonnegative. The second condition is the key to understanding how contracting frictions in period 2 can prevent trade from occurring in period 1, even when there is positive surplus. If μ and λ are sufficiently low, then trade may not occur.

2.3. **Predictions.** First note that, if $\mu = \lambda = 1$ (this is equivalent to perfect enforcement), then for each subsidy case, condition (2) is irrelevant—it is automatically satisfied whenever

[9]In addition, our model predictions hold if these are fixed costs in levels, rather than a proportional decrease.

condition (1) is satisfied. In other words, in the absence of period 2 contracting frictions, trade will occur anytime there positive surplus from trade. This is the essence of the Coase Theorem benchmark result in the absence of frictions. In addition, the buyer's wealth on hand is irrelevant to whether trade occurs. This is because, as long as there are gains from trade, the seller can recover funds from the surplus generated in period 2.

In Proposition 1, the second condition encapsulates the core difference between the two subsidy cases. In the Seller subsidy case, the seller knows he will receive s for sure in period 2. In the Buyer subsidy case, the seller must trust the buyer to transfer s to him. If any agreement in period 1 is perfectly enforced in period 2 (i.e. $\lambda = 1$), then the second condition in both cases is the same; consequently, there will be no difference in the level of trade in the two subsidy cases. In contrast, if the seller cannot perfectly trust the buyer to deliver the subsidy to him in period 2 (i.e. $\lambda < 1$), then condition (2) will be weakly less likely to be satisfied in the Buyer subsidy case relative to the Seller subsidy case; consequently, trade will be weakly lower in the Buyer subsidy case relative to the Seller subsidy case. This is the key prediction of our model.

In addition, note that our test will only have power to detect period 2 contracting failures if w is sufficiently small. If the buyer has sufficient up-front wealth so that $w \geq c$, then condition (2) will be satisfied in both cases and the level of trade will be the same under both subsidy cases—even if $\lambda = 0$. Intuitively, this is because the buyer can satisfy the seller's participation constraint by giving him adequate compensation up front in period 1. This means trade can occur even if period 2 frictions are severe, because the parties can do everything ex ante, making period 2 frictions irrelevant.

This observation indicates a potentially important role for the buyer's wealth in whether trade occurs. First, note that increases in the buyer's wealth will make constraint (2) more likely to be satisfied in both cases—weakly increasing the probability of trade regardless of who receives the subsidy. In addition, the conditions in Proposition 1 suggest that the buyer's wealth may prove to be more binding in the Buyer subsidy case. Specifically, in the Buyer subsidy case, if λ and μ are small (i.e. close to zero), then trade will not take place unless the buyer has enough liquid wealth in period 1 to compensate the seller for c. In contrast, under the Seller subsidy case, a more modest wealth level is sufficient to enable trade, since the seller is assured to receive s in period 2. However, our model does not make predictions on which subsidy case will be more responsive to increases in w.[10]

2.4. **Discussion.** Conceptually, the introduction of the subsidy gives us a lever to manipulate ex post enforcement levels—through our ability to vary which party receives the

[10]This depends on the distribution of w in the buyer population, and its level in relation to c, λ, and μr. Without taking a stance on these parameters, our model does not make clear predictions on whether increases in w will have a bigger impact on the level of trade in the Buyer subsidy or Seller subsidy case. In addition, the model implicitly assumes that the seller has the liquidity to cover the cost of c up front; if the buyer and seller cannot cover the costs of irrigation between them, then then trade may not occur in either the Seller or Buyer subsidy cases, and we may not be able to detect enforcement failures even if they exist.

subsidy. In the experiment, we randomize the recipient of the ex post subsidy and test whether this affects the probability of ex ante trade.

Could a mechanism other than costs of future funds recovery generate a difference in trade if the subsidy is delivered to the seller instead of the buyer? In the experiment, one potential consideration is that sellers and buyers have different beliefs about whether we will return to deliver the subsidy payment. Suppose sellers are more likely to believe we will return than buyers on average. In this case, the sellers in *both* treatment groups would be more likely to believe that we will deliver the money as promised; if there are no enforcement constraints, then, based on their respective beliefs, buyers and sellers will agree to an ex post division of the subsidy, and outcomes should look no different in the two groups.

Similarly, if the buyer has a higher alternate use of funds (e.g. due to a negative shock), this will be the case in both treatment groups; as long as the seller trusts that the ex post division will be as promised, in whose hands the money arrives should not matter. Alternately, if bargaining power is affected by who receives the funds, this may affect the ex post division of surplus but should not affect whether trade occurs. In short, if there are gains from trade in the case of the Seller-subsidy treatment, then these gains exist in the Buyer-subsidy treatment; outcomes on the amount of trade should look the same as long as the seller trusts the buyer will split the subsidy as promised.

Finally, the model above assumes perfect and symmetric information among the buyer and seller—specfically, that b is known with certainty. However, more general forms of incomplete contracting (aside from enforcement constraints) could prevent trade from happening even when $b > c$. In the experiment, such problems will be common to the Buyer and Seller subsidy groups. By introducing an external subsidy amount s and explicitly informing both parties about s, we are able to ensure that information about s is symmetric. In general, one could write a more complicated incomplete contracting model where altering who receives the subsidy leads to differential trade. Consequently, we recognize that the most defensible interpretation of our design is a test for contracting failures or transaction costs. We view our intervention as being most consistent inducing a change in the certainty or cost of future funds recovery, and in the exposition, we will use this language in what follows for concreteness.

3. Experiment Design

3.1. **Context: Groundwater Markets.** We conduct our test in the context of spot markets for groundwater in the central/eastern region of the state of Uttar Pradesh, India. In this area, groundwater is the predominant source of irrigation water for agriculture. The fixed cost of sinking a borewell and purchasing an engine to pump the water out of the ground is fairly large. Borewells are therefore typically owned by wealthier farmers in a village.

Farmers who do not have their own well can purchase irrigation from a well-owner on a neighboring plot of land. There are extremely active spot markets for groundwater in the region. Buyers typically rent another farmer's borewell and engine at an hourly rate. 99% of the water transactions in our baseline survey sample were these hourly spot contracts (in contrast with season-long irrigation contracts which are prevalent in other parts of India). The hourly rate includes the variable cost of diesel, which is used to power the motor and accounts for about 50% of the hourly market price. Another implicit variable cost is the seller's time: irrigating another farmer's land takes the bulk of a seller's day, because sellers remain with the engine while it is running and due to fixed set up costs.

While irrigation purchases can happen all year, the primary irrigation season is from May-June—the hottest months of the year. Farmers who choose to grow crops during this time of year, particularly sugarcane, must irrigate to prevent crops from drying out.

Water is transported via cheap plastic hoses that can be attached to the well and run to the desired plot of land. Because there is loss in water from transport over long distances, farmers typically only purchase water from someone on a nearby plot of land. Most farmers have access to 1-5 potential sellers, with a mean of 3 in our sample. Buyers typically purchase water from a neighboring farmer multiple times each year. This is therefore a setting with a high degree of repeated interactions: buyers and sellers are neighbors and will be for their entire lives (given limited mobility and extremely low levels of land sales).

Figure 2 indicates that 99% of water buyers perceive the net returns to an additional irrigation on agricultural profits to be positive. In the baseline survey, we asked water buyers how much their crop revenues would increase if they irrigated one more time for each crop they grow, and how much one additional irrigation could cost for each of these crops. Most perceive the magnitude of the returns to be fairly high, with a mean net return of 123% averaged across all crops. These perceived returns are even higher if we focus on only the plot where the farmer perceives the return to be highest, with half of respondents citing a net return above 100% (Appendix Figure A1). During focus groups, smallholder farmers claim they do not irrigate as much as they'd like due to liquidity constraints.

Water sellers typically demand payment for irrigation at the time of sale or shortly after. At baseline, water buyers described purchase arrangements with each seller from whom they had bought irrigation in the past year, for a total of 337 buyer-seller transactions. In all instances except two (99.4% of instances), the buyer said he paid within a week of irrigation. In only one instance did the buyer wait until harvest (when returns to irrigation would be realized) to pay a seller (Figure 3, Panel A).[11] Consistent with these self-reports, in

[11]In contrast, there is some evidence that wealthier farmers may be more likely to delay payment to sellers. Because landholdings are fragmented, those who own a borewell on one plot may purchase water from a seller for another plot. Among the water sellers in our sample who also engaged in water purchases, we observe a somewhat higher incidence of waiting for long periods or after harvest to repay their sellers. This could reflect credit-worthiness, social ties among sellers, or some other factor. We unfortunately do not have enough variation to explore this heterogeneity more carefully in the experiment.

transaction-level data from the control reference group in our experiment, 69% of payments were due within a day of irrigation, and 95% of payments were due within 5 days (Figure 3, Panel B). There were no transactions where a seller allowed the buyer to defer payment for more than 10 days after the time of irrigation.[12] This indicates that, although the fixed cost of the borewell has been sunk and water sellers tend to be among the wealthier residents in a village, they are unlikely to allow their neighboring water buyers to wait to pay them. Of course, the patterns in Figures 2 and 3 do not necessarily imply any inefficiency. For example, if there is high demand for irrigation, the market may clear at terms that are favorable to sellers, such as up front payments.

In this paper, our goal is not to explain the equilibrium in the irrigation market. Rather, the stylized facts above—the high perceived returns to irrigation, limited intertemporal contracting, and repeat interactions among experienced agents who are neighbors—suggest that this market provides a good backdrop for our empirical tests.

3.2. **Pair Construction.** We identified potential water buyer-seller pairs in 21 villages. In each village, we constructed a census of cultivators. We identified potential "water buyers" as farmers who cultivated a plot of land without a well on it, and randomly picked a subset of these in each village. For each of these chosen water buyers, we identified all neighbors with a borewell and pump engine who were close enough to potentially sell water to that buyer; we randomly picked one of these potential sellers. Any given household could only be a part of one pair. Our sample is comprised of 407 pairs.[13]

3.3. **Treatments and Randomization.** To implement the test laid out in our model, we introduce a subsidy that is awarded to a buyer-seller pair if they trade with each other. The pair receives the lump sum subsidy each time the buyer purchases irrigation from the seller over a 3-month period. The participants were told that the subsidy payment would be delivered *after* the end of the irrigation season, in July. This meant that if the pair did want to irrigate, they would need to come up with the funds to do so among themselves during the irrigation season. The subsidy was substantial in size—constituting about 50% of the cost of a typical irrigation.

Each pair was randomized into one of the following treatment groups:

(1) Seller Subsidy: Subsidy payment delivered into the hands of the water seller.
(2) Buyer Subsidy: Subsidy payment delivered into the hands of the water buyer.

[12] During focus groups, sellers claim that payment recovery is easier in the case of short term trade credit (within a week) relative to longer horizons—for example, because the state of the world is more verifiable (the buyer cannot claim a change in his household situation). This could be a potential reason why sellers provide shorter term trade credit but not longer-term credit.

[13] 417 unique buyer-seller pairs took part in the experiment. However, we did not collect data from 10 of these pairs in our first endline survey. For consistency across tables, the analysis is presented for the 407 pairs (i.e. 98% of the sample) for which the core experiment data (weekly checks and survey at the end of the irrigation season) was collected. The results are essentially unchanged if we re-run the analysis for all 417 pairs for treatment effects on irrigation levels.

(3) Control: No subsidy offered.

Before the start of the irrigation season, both members of the pair were informed together about the details of the subsidy offer: the amount, timing of payment delivery, and to whom it would be delivered. We stratified randomization by village, with 40% of pairs within a village assigned to each of the two subsidy conditions and 20% of pairs assigned to the Control group.

Given the magnitude of the subsidy, the gains from trade are substantially higher in Groups 1 and 2 than in Group 3. We therefore would expect that the subsidy groups (Groups 1 and 2) would irrigate more than Group 3.

Our primary interest is in comparing the level of trade between the two subsidy groups. Treatments 1 and 2 mirror the two cases in the model. The seller bears the cost of irrigating (in terms of diesel, his time, and possible depreciation of the engine) at the time of irrigation. The buyer can compensate the seller for these costs at the time of trade, or potentially defer some part of the payment until a later date (delivery of our subsidy 3 months later, or after harvest). Note that the total surplus from trade, timing of events, information available to each party, and liquidity available at the time of trade is exactly the same in both Groups 1 and 2. The only difference is whether the seller is assured of receiving the subsidy payment directly, or whether it goes to the buyer—creating a potential recovery issue. If the buyer and seller can agree at the time of trade on how to divide the subsidy, and expect that both parties will follow through on this without reneging when the subsidy arrives, then there should be no difference in the amount of trade during the irrigation season between Groups 1 and 2. However, if the seller perceives recovering money at a later date to be costly—for example, because there is a chance that the buyer will renege or making visits to collect money from a neighbor violates social norms—then the level of ex ante trade during the irrigation season will be higher under the Seller Subsidy than under the Buyer Subsidy.

3.4. **Timeline and Protocols.** Figure 4 summarizes the experiment timeline. Our experiment census was conducted in the year before the intervention ocurred, and was the sample frame from which we selected participants and plots.[14] We approached buyer-seller pairs in early March, after planting decisions for the season have been made. At this time, for each pair, we conducted a meeting that included the buyer, the seller, the elected village head (pradhan), and one of our field staff. For pairs in the subsidy groups, the field staff member explained the rules of the subsidy offer, as described above. The treatment was implemented using a small manipulation in the script: during the meeting, at the end of the explanation, the pair was told we would return in July and "the subsidy money will be delivered to <<participant's>> house", with the name of the buyer or seller filled into

[14] Because the participants and sample plots were chosen well in advance of the experimental offers due to logistical reasons, we ran the risk that some plots would be fallow (a decision made before any participants knew about the experiment). This did indeed happen for some plots. These plots are included in the analysis; as one would expect, results get stronger if we exclude them.

the script at the end, based on the pair's treatment assignment. All other aspects of the interaction remained the same. For pairs in the control group, the field staff simply reiterated that the buyer and seller could potentially trade with each other during the upcoming irrigation season.

The purpose of having the village leader present at each sit-down was to build confidence that we would indeed return three months later with the subsidy payment as promised. We also built trust with participants in two additional ways. First, we had conducted baseline surveys several months earlier in the villages where the experiment was conducted, and households were paid for their participation. Many participants were therefore familiar with us and had received money from us in the past. Second, our staff visited the buyer and seller every week during the irrigation season, making them a regular and familiar presence in the village while the experiment was being conducted.

Any irrigations conducted between April and June were eligible to count for the subsidy payments. While the actual irrigation season is in May-June (the hottest months of the year), we included April in the subsidy window, since this is when irrigations could potentially begin. Buyers and sellers could irrigate as many times as they wanted during this period, with a lump sum subsidy amount s earned for each irrigation instance. The participants were told that the total earned subsidy payments would be delivered in cash in the beginning of July (to the buyer or seller, based on the pair's treatment assignment). Harvest for sugarcane—the predominant cash crop in the area, and the crop for which irrigation is most frequently purchased during this time—occurs starting in October and continues until the following January. The payoff to irrigation, in terms of crop revenue, would therefore be realized 3-6 months after the end of the irrigation season.

3.5. **Data and Descriptive Statistics.** To accurately measure trade, we surveyed each pair weekly during the experiment. Every week, our surveyors visited each buyer and seller separately to ask them if they irrigated; if they both reported they had, the staff walked to the buyer's plot to verify irrigation by checking the soil moisture.[15] They also collected information on the number of hours of irrigation purchased, the price charged, and the date at which payment was made or was expected to be made. A year after the intervention, we performed an endline survey through which we collected data on crop yields and retrospective information on the amount of irrigation that had been purchased from other potential sellers (aside from the paired seller).

[15]This allows us to verify that the buyer's plot was indeed irrigated. The fact that we see meaningful increases in yields also supports the view that actual irrigation increased. Of course, even if the soil is moist, the irrigation could have been purchased from someone other than the paired seller. Even if the buyer and seller colluded to lie to us about who performed the irrigation, our basic contracting test still holds: if there is less irrigation under the Buyer subsidy than the Seller subsidy, this means the Buyer subsidy group was unable to replicate the arrangement that was possible under the Seller subsidy, and therefore left money on the table.

Finally, we have two sets of basic demographic and baseline measures. The first set was collected from the full experiment sample. In addition, we conducted a fuller, more detailed baseline survey for a subset of the households in the experiment. This data is from an extensive survey that was conducted in many villages in the area a year before the intervention, as part of a broader project. 61% of the buyers and sellers in the experiment were part of this full baseline survey sample.[16]

Table 1 provides summary statistics and balance checks on baseline covariates. Buyers are active participants in the water market: 94% of buyers purchased irrigation for one of their plots at some point last year. In addition, 62% of buyers had purchased irrigation from their paired seller in the past. The borewell owners from whom buyers can puchase irrigation is limited to a handful of their neighbors (for an average 2.99 possible sellers); in 14% of cases the paired seller is only possible seller.

Because pair members work on neighboring small plots of land and reside in the same village communities, in general, buyers and sellers tend to know each other personally. In 59% of pairs, both the seller and buyer have each visited inside of the other's home in the past—a sign of social closeness in a setting where caste hierarchies can limit who visits inside one's home. In about half the pairs (49%), the seller and buyer are the same jati, or subcaste. At the same time, there are wealth differences between buyers and sellers, as indicated by average asset holdings and consumption; in 73.5% of pairs, the water seller has higher annual consumption than the buyer.

Table 1 indicates that treatment assignment was reasonably balanced between the two subsidy groups (Col. 7). Aside from relative consumption levels (where the difference between the subsidy groups has a p-value of 0.092), the difference in covariate means between the two groups is insigificant—albeit with some magnitudes that are meaningful in size.[17] As we verify below, our results are robust to a variety of control strategies and specifications.

4. Results I: Trade and Output

4.1. **Take-up of the Subsidy.** Recall that the subsidy covers 50% of the market costs of the typical irrigation. Despite this incentive, 62% of buyer-seller pairs in the Subsidy groups never trade with each other (Appendix Figure A2). This suggests that the subsidy offer was not strong enough to enable trade in most pairs, despite baseline beliefs among water buyers that the returns to irrigation are high. This could be due to liquidity constraints or technological match-specificity among buyers and sellers. Alternately, some of this may reflect another class of contracting barriers. For example, if the buyer and seller were

[16] In the empirical analysis, for those controls where we do not have data for all participants (i.e. the 39% of participants who did not participate in the extensive baseline survey), we code those values as zeros and add dummies to indicate missing baseline data in the regressions.

[17] Because the control group was comprised of only 20% of the sample, the difference between each subsidy group and the control group (Cols. 3-6) has some additional coefficients with significant differences, as may be expected due to chance.

different religions (i.e. Hindu and Muslim), the pair never traded and therefore did not use the subsidy offer. While not the focus of our test, which is designed to examine intertemporal breakdowns, we discuss this further in Section 5.3 below.

4.2. **Effects on Trade (Irrigation).** To examine treat effects on trade, we begin by plotting the simple average of hours of irrigation purchased in each week of the experiment, separately for each treatment group (Figure 5). The figure shows the total unconditional mean for each group, including zeros for pairs that did not trade. The amount of irrigation picks up for all groups after week 5—denoting the start of the irrigation season in May, when extreme heat begins. Irrigation purchases taper off by week 14, as the monsoon onset begins in early July.

In accordance with our predictions, Figure 5 shows: (i) Both Subsidy groups irrigate more than the Control group, and (ii) the Seller subsidy group irrigates more than the Buyer subsidy group. This ranking is robust—holding across weeks on average. This supports the view that, when the seller anticipated receiving the subsidy directly (rather than it going to the buyer), trade was more likely to occur during the 3 months before subsidy arrival.

To estimate treatment effects in a regression framework, we estimate Intent to Treat regressions throughout the paper:

$$y_{ij,t} = \beta_0 + \beta_1 SellerSubsidy_{ij} + \beta_2 BuyerSubsidy_{ij} + \delta_v + \mathbf{X'_{ij}}\theta + \varepsilon_{ij,t}, \tag{4.1}$$

where $y_{ij,t}$ is the amount of irrigation between buyer i and seller j in week t. $SellerSubsidy_{ij}$ and $BuyerSubsidy_{ij}$ are dummies for whether buyer-seller pair ij was assigned to the Seller subsidy group or Buyer subsidy group, respectively. The omitted category in the regression is assignment to the Control group. The δ_v is a vector of village fixed effects, and $\mathbf{X'_{ij}}\theta$ is a vector of baseline covariate controls.[18] Under the null of perfect enforcement, we would expect $\beta_1 = \beta_2$. However, under enforcement costs, we would expect $\beta_1 > \beta_2$. In addition, since the subsidy increases the gains from trade relative to the Control, we expect $\beta_1 > 0$ and $\beta_2 > 0$ (regardless of whether there are contracting frictions).

Table 2 shows estimated treatment effects. Panel A examines trade during the irrigation season. Col. (1) provides OLS estimates.[19] Because of the large percentage of zero values in the hours of irrigation, we estimate a tobit model in Col. (2) of Table 2 and report marginal effects. The pairs in each subsidy group trade more than the pairs in the Control group. In addition, the Seller subsidy group trades substantially more than the Buyer subsidy group: the difference in the hours of irrigation purchased is 0.649 hours per week in the main irrigation season (p-value of 0.023). In Panel B, the pattern of results for the full subsidy

[18]We control for the buyer's baseline crops and proxies for the major sources of heterogeneity that we explore below and which would be expected to mediate trade: wealth, caste distance, and market power (whether the paired seller is the buyer's only potential seller). We show that estimates are similar under a variety of specification checks below. Note that the experiment was run before the AEA RCT registry was established.

[19]Appendix Table A1 shows robustness to alternate specifications.

period is similar: the Seller subsidy group irrigates 0.55 more hours per week on average. This magnitude corresponds to 58% of the Buyer subsidy group mean.

In Col. (3), we restrict analysis to pairs that have traded in the past, and consequently have a prior market relationship.[20] The effects hold strongly for this group, indicating that such pairs are not able to overcome the contracting problem. Rather, among pairs that have not traded in the past (Col. 4), there is no take-up of the subsidy offer—with both subsidy groups having similar levels of trade as the control group. This indicates that if a farmer has not already engaged in trade with one of his neighbors, then the 50% subsidy is not sufficient incentive to get them to trade.

Finally, in Cols. (5)-(8), we repeat the analysis on the extensive margin, with a binary dependent variable that equals 1 if the pair irrigated in that week. The same pattern of results holds.[21] Overall, these results verify our main prediction: the amount of trade falls substantially if the parties expect the subsidy to be delivered to the buyer rather than the seller.

The above results are for trade within the pair. In Table 3, we examine whether trade within the pair crowds out the buyer's purchases of irrigations from other potential sellers. We find scant evidence of crowd-out: buyers in the subsidy groups are not more likely to switch away from their other sellers (Col. 1-4) and do not reduce the number of irrigations they purchase from others (Cols. 5-6). Consequently, in the subsidy groups, the estimated number of irrigation hours from other sellers is not different from the control group or each other—magnitudes are small and statistically insignificant (Col. 7). Col. (9) estimates effects on total irrigation purchased across all sellers (the paired seller + other sellers). The Seller subsidy leads to a net increase in the total amount of water irrigated on the buyer's plot: 4.6 hours (38%) more than in the Buyer subsidy group.[22] Col. (10) re-runs the analysis among pairs who had traded in the past.

While the subsidy greatly reduces the effective cost of trade, buyers do not substitute from other sellers to this cheaper source of irrigation. Some substitution may potentially be expected if buyers were unconstrained in irrigation purchases at baseline (see Banerjee and Duflo 2014). Rather, the subsidy—particularly the Seller subsidy—leads to a net expansion in the amount of irrigation. In Appendix Table A2, which is discussed in more detail below, we verify that the subsidy also led to a net increase in the overall amount of irrigation sold by sellers (Cols. 9-10).

[20]It would be preferable to have more detailed heterogeneity in frequency of previous trade, but we unfortunately do not have such data.

[21]Because pairs received the subsidy for each irrigation (regardless of number of hours), one concern would be that participants reduced the number of hours per irrigation—stretching the same number of hours across more irrigations. Cols. (1)-(2) indicate that the underlying amount of water that was transferred to the buyer's plot was higher in total. In addition, conditional on irrigation, the Seller subsidy group actually irrigated more hours per irrigation than the Buyer subsidy group (p-value 0.087, results available on request).

[22]The Tobit marginal effects estimates are similar, due to the reduced number of zeros, with similar levels of statistical significance in the difference between the two subsidy groups (available upon request).

4.3. **Effects on Output (Crop Yields).** Do the effects on irrigation translate into meaningful impacts for farmers? Because treatment offers were made after planting decisions were made, there is no scope for the subsidy offer to affect crop choice; we verify this in Appendix Table A3. Consequently, to examine impacts on output, we take crop type as given and examine effects on yields.

Table 4 examines the reduced form effects of the subsidy treatments on crop yields. Crop yields is a composite index, standardized using the means and standard deviations of yields in the Control group. [23] Relative to the Control group, buyers' crop yields increase by 0.411 standard deviations in the Seller-subsidy group (Col. 1, significant at 5%); the point estimate for the Buyer subsidy condition is 0.076 but not significantly different from yields among the Control group. Consequently, relative to Buyer-subsidy, the Seller-subsidy increases crop yields by 0.335 standard deviations (p-value 0.047). Col. (2) estimates effects for pairs that have a prior market relationship, and finds a similar pattern.

Overall, relative to Buyer-subsidy, the buyers in Seller-subsidy have 9.7% higher crop revenue (Col. 3), though this is more noisily estimated. While this magnitude is large, it constitutes an upper bound on the impact on the buyer's take-home profits. Since we did not collect data on the use of other inputs (e.g. labor, fertilizer), which may also have increased along with irrigation, we cannot compute treatment effects on profits or the returns to irrigation directly. However, we can use the baseline survey data to obtain one suggestive benchmark: in the cross-section at baseline, a 0.335 standard deviation increase in the yields index corresponds to an approximately 3% increase in self-reported crop profits.

As a whole, the results in Tables 2 and 4 suggest that the magnitude of effects on irrigation is economically meaningful. Buyers appear better off—at least by these measures—when the subsidy payment is delivered to the seller rather than to themselves.

5. Results II: Contractual Structure and Determinants of Trade

5.1. Contractual Terms.

5.1.1. *Ex-post Transfers: Sharing of Subsidy Payments.* Did pairs who traded share the subsidy after it was delivered? When the subsidy payments were delivered, we asked each pair member separately how they intended to divide the subsidy payments. There are two stylized features of the responses.

[23]Note that, due to an oversight during endline data collection, we did not collect endline yields data from all pairs. This was an administrative error in survey collection, leading to 10% of missing yields observations. These respondents answered other modules in the endline survey, and so we do not view this as reflecting selective attrition. To check this, Appendix Table A4 shows regressions of a missing yields dummy on each covariate in the balance table; as expected, there is no evidence of correlation with observables (Panels A-C). In panel D, we verify there are not significant differences in the observation rate across treatment groups; e.g., the Seller subsidy group is 1 percentage point less likely to have a yields observation than the Control group.

First, there was a remarkable degree of ex ante agreement among the buyer and seller on how the subsidy would be divided. 124 pairs received a positive subsidy payment. In 90% of these pairs, the buyer and seller gave us the same exact answer about the amount that would go to each party. In the remaining 10% of cases, one member of the pair said the share amount was "Undecided". In cases where sharing would occur, both the buyer and seller gave the same specific amount that would be transferred to the person who didn't directly receive the subsidy, with amounts ranging from Rs. 200 to Rs. 5,250. This suggests that almost all the subsidy pairs that traded had either explicitly discussed specific amounts beforehand, or had a common understanding of what would happen at the end of the season when the subsidy was delivered.

Second, overall, only 18% of pairs decided that they would share the subsidy payment after it was delivered; in the remaining pairs, they stated that the subsidy recipient would keep the entire subsidy amount (Appendix Figure A3). Among the pairs that did share the subsidy payment, the proportion of the subsidy that was transferred by the recipient to the opposite party ranged from 14% to 75%, with a mean of 40%.[24]

This pattern is consistent with the idea that the individuals in our sample are reluctant to enter into contracts that require future exchange of funds. If there are costs or barriers to funds recovery, parties will ex ante decide not enter into contracts that involve ex post transfers.

In addition, if most individuals expected the subsidy recipient to keep the payment, this would pose a larger hurdle for the Buyer subsidy treatment than the Seller subsidy treatment. This is because the Seller is the first mover (i.e. he must pay the cost of irrigation up front if irrigation is to happen). Consequently, his expectation of how much money he will get in the future would constrain whether trade occurs; this participation constraint is more likely to be satisfied in the Seller subsidy case than the Buyer subsidy case.

5.1.2. *Ex-ante Transfers: Price changes.* If the subsidy recipient cannot credibly commit to transferring funds ex post, he could make an ex ante transfer at time of sale—for example, via a change in up front price paid. Such transfers could be used to secure the other party's participation in trade. We examine features of contract terms to check for such transfers. Note that such analysis is necessarily suggestive: the model delivers clear predictions on the level of ex ante trade. However, the effects on contract terms reflects division of the surplus, for which clear predictions cannot be formed without taking a strong stance on the bargaining technology.

[24]In the 10% of pairs where the buyer and seller did not agree on their answer, the subsidy recipient had a firm answer but the opposite party said "Undecided". In all these cases except one, the subsidy recipient said they would keep the full amount. In Appendix Figure A3, for these 10% of cases, we code the sharing rule as the one that was given by the subsidy recipient. In addition, we did not go back to verify if the subsidy was actually divided the way respondents said it would be. Consequently, these numbers can likely be interpreted as an upper bound on how much sharing there was.

Since we only observe prices when trade actually occurrs, this complicates the interpretation of any potential price effects. For example, if the marginal trade induced by the subsidy is less desirable than those in the control group, then the expected price required to satisfy the seller's participation constraints under the Subsidy treatments may need to be higher. Similarly, if sellers' cost to irrigation is convex (e.g. due to capacity or time constraints), then the expansion in irrigation sales under the Seller subsidy may require a higher price level. Such factors would make it more difficult to observe price reductions offered by sellers under the Seller subsidy. Similarly, liquidity constraints on the part of buyers, if they prevent irrigation sales under the Buyer subsidy, would make it more difficult to observe underlying price changes that would have occured under Buyer subsidy irrigation sales.

With these caveats in mind, Table 5 examines the difference in prices charged for irrigation among the treatment groups.[25] On average, buyers pay Rs. 20/hour less in the Seller subsidy group, relative to the Buyer subsidy group (Col. 1, p-value 0.078). Overall, they are 6.9 percentage points (82%) more likely to receive a discounted (Col 2, p-value 0.015). These findings are consistent with sellers being relatively more willing to lower prices up front to induce the buyer to trade when they know they will be receiving the subsidy. When a price reduction does occur under the Seller subsidy condition, it corresponds to 39% of the market price of the average irrigation. While we observe differences between the two subsidy groups, estimates are too noisy to distinguish price differences relative to the Control group.

5.1.3. *Trade credit.* When the seller expects to receive the subsidy directly, he is also more likely to extend short term trade credit to buyers. Figure 6 shows the cummulative distributions of number of days before payment is due, conditional on an irrigation. The distribution for the Seller subsidy group is shifted to the right of the Buyer subsidy group. However, note that the x-axis of Figure 6 makes it clear that sellers do not allow buyers to wait until subsidy delivery (or harvest) to repay them. In 90% of sales, payment is due within 3 days of irrigation. The maximum duration of deferred payments is within 20 days of irrigation.[26]

[25]There is usually a standard going hourly rental price for a borewell in each village. In our sample, this going price was either Rs. 70/hour or Rs. 60/hour. We use the modal hourly price in the Control group in each village to determine the going village rate. Very few observations deviate from this modal price. The market price for an irrigation is the hourly rate*number of hours irrigated. The dependent variable in Col. (1) is the Amount charged - Market price.

[26]These numbers show the payment due date on the date of the irrigation transaction. Unfortunately, we do not have data on whether each of these deferred payments was made on time. Consequently, this is likely a lower bound on how long sellers waited to receive payments. In addition, in seven instances, the buyer had not yet paid the seller when we arrived for the weekly survey, but respondents did not give an expected payment date. We code these as being due within a week of survey date, but the results are robust to alternate codings. In addition, in the regression results below, we show effects on a binary indicator for any deferred payment, which does not depend on a coding choice for these 7 observations.

Such small payment extensions are presumably valuable to buyers because they allow irrigation to be scheduled within the seller's (sometimes inflexible) preferred irrigation date, while allowing buyers some buffer to come up with the funds—e.g., by working in the casual labor market. Why might sellers be willing to allow short-term deferrals of payment (e.g. a few days), but be unwilling to wait 3 months (when the subsidy is delivered) to collect funds from buyers? While we do not have direct evidence on this, during focus groups, sellers said that payment recovery becomes more difficult and uncertain if the time horizon is longer. They claimed this is because the buyer has greater scope to claim a change in the state of the world (e.g. a household health shock that depleted savings) or allows for increased monitoring (e.g. the seller can verify that the buyer worked in the casual labor market to generate the funds).[27] This view is consistent with the survey evidence presented in Section 6 below, in which sellers state that multiple repeated attempts are needed to collect funds under longer time horizons.

Table 6 presents regression results on deferred payments. It confirms that sellers extend more trade credit to buyers in the Seller subsidy condition relative to the Buyer subsidy condition. On the extensive margin, they are 2 percentage points more likely to offer any credit—a 39% increase relative to the Buyer subsidy mean.

5.1.4. *Transfers - Summary.* Appendix Table A2 provides evidence that in the Seller subsidy condition, sellers reduce the number of other buyers they trade with (Cols. 1-5). This is only a partial crowdout. The total amount of irrigation performed by sellers across all their buyers still increases (Cols. 9-10). Although groundwater was fairly abundant in this area, this partial crowd-out could reflect time constraints, since irrigating a farmer's land takes most of the workday (see Section 3.1above). Sellers are willing to (partially) switch away from their other buyers in the Seller subsidy condition, but not in the Buyer subsidy condition.

Together, the above findings suggest that sellers engage in a variety costly concessions or actions to take enable trade when they are assured the subsidy payment. The price reductions and trade credit increase the likelihood that the buyer will participate in trade—potentially both by reducing the liquidity he needs to have on hand at the time of irrigation, and also by making trade more attractive. While sellers appear willing to make such up front concessions under the Seller subsidy, we see little evidence they are willing to do so under the Buyer subsidy.

5.2. **Heterogeneity: Mediating Effects of Wealth.** Proposition 1, condition (2) requires that under the Buyer subsidy case: $w + \mu r + \lambda s \geq c$, and under the Seller subsidy case: $w + \mu r + s \geq c$. This condition highlights the role of the buyer's wealth in enabling

[27]In addition, some dynamic moral hazard models would predict that it is optimal to secure payment for one irrigation before irrigating again; this prevents the accumulation of a large sum of debt, which would lead to a larger temptation to renege.

trade. In the absence of enforcement costs, w will play no differential role in enabling trade: as long as condition (1) is satisfied, it will not matter if funds exchanged up front at the time of irrigation or in period 2 (i.e. when the subsidy is delivered). In contrast, under enforcement costs, the buyer's wealth is important because it can enable him to make ex ante transfers at the time of irrigation—for example, by making it more likely that the buyer can pay the full amount for the irrigation up front, making a price discount or trade credit unnecessary. An increase in the buyer's cash on hand loosens constraint (2) under both the Buyer subsidy and Seller subsidy case, weakly increasing the probability of trade under both treatments. In addition, if the cost of future funds recovery are sufficiently high (i.e. μ and λ are sufficiently small), then trade will be impossible under the Buyer subsidy treatment unless the buyer has sufficient wealth upfront. In contrast, this need not be true in the Seller subsidy case, since the seller is assured s with certainty.[28]

Table 7 examines treatment effects on irrigation by baseline wealth. We examine two baseline proxies for wealth: total asset value and consumption in the previous year. The coefficients in Panel A show estimated effects for the case when the buyer's wealth is low (below median). For each of the wealth proxies, trade under the Buyer subsidy is not significantly different from trade under Control—both in terms of the hours of irrigation (Cols. 1-2) and the probability of irrigation (Cols. 3-4). This suggests that when they are poor, buyers have a difficult time taking advantage of the Buyer subsidy. In contrast, when buyers are wealthier (Panel B), there is robust take-up of the Buyer subsidy treatment, with significant differences relative to the Control group.

Under the Seller subsidy treatment, there are strong treatment effects on trade even when the buyer is not wealthy (Panel A). While increases in the buyer's wealth leads to increases in trade under the Buyer subsidy across both the assets and consumption proxies, the pattern for the Seller subsidy is less consistent.

Consequently, in the Buyer subsidy treatment, the pairs cannot even take advantage of the subsidy unless the buyer is likely to have cash on hand. In contrast, we see no evidence that increases in the Seller's wealth lead to higher trade under the Buyer subsidy treatment (results available on request). Overall, these findings illustrate a way in which the enforcement problem can exacerbate the problem of underdevelopment: it is less likely to be relevant when parties are wealthier.

5.3. **Heterogeneity: Potential Correlates of Informal Enforcement Mechanisms.** A large theoretical literature in economics establishes ways in which relational contracting

[28]We should note that our model does not produce crisp predictions on whether increases in w will have a bigger impact on trade under the Buyer subsidy vs. the Seller subsidy. This will depend on the distributions of paramter values in condition (2). For example, even if $\lambda = \mu = 0$, this prediction will still depend on the relative size of s versus c, and the underlying distribution of w across buyers. Consequently, we do not make predictions here about whether the heterogeneous effect of wealth increases should be larger under the Buyer or Seller subsidy. Rather, we focus on the clear qualitative prediction that, if λ and μ are small, there will be low levels of trade under the Buyer subsidy unless w is sufficiently large.

can enable agents to achieve the first best level of contracting, despite lack of formal enforcement mechanisms. In Table 8, we undertake an exploratory analysis to examine potential correlates of informal enforcement mechanisms. For each covariate listed at the top of the table, we report estimated treatment effects when the covariate equals 0 in Panel A and when the covariate equals 1 in Panel B.

We begin with proxies for market power. Among 13.8% of pairs, the buyer has no other potential sellers near him aside from his paired seller. In such pairs, because the buyer cannot turn to someone else for irrigation, the seller has an especially powerful trigger strategy—potentially making it more likely that buyers will follow through on promises, and therefore mitigating the costs of future funds recovery. In Col. (1), we examine heterogeneity in treatment effects along this dimension. When buyers have multiple potential sellers (Panel A), we see large and statistically significant differences in irrigation between the Seller subsidy and Buyer subsidy. However, if the paired seller is the buyer's only potential seller (Panel B), the Buyer subsidy coefficient is higher and we no longer see significant differences between the two subsidy groups. In Col. (2), we see a similar pattern of results when the buyer is the seller's only potential buyer. The patterns in Cols. (1)-(2)—while broadly consistent with the idea that market power increases pairs' ability to take advantage of the Buyer subsidy, thereby bringing irrigation under the two subsidy groups closer together—are noisy and therefore only suggestive. We are under-powered for this analysis because these covariates equal zero for most pairs. Linear specifications or more balanced cuts of the data suggest that the contracting benefits only hold in the extreme case when one of the members of the pair is the only potential trading partner. Consequently, even if this channel enables pairs to overcome contracting frictions, in practice it holds for a modest fraction of buyers and sellers.

In the remaining columns, we examine correlates of social distance. If the buyer and seller have strong social ties, this could improve contracting, for example, through greater levels of underlying trust or enforcement through the social or caste network. If such ties mitigate the costs of future funds recovery, then we would expect them to increase trade among Buyer subsidy pairs, reducing the difference between the Buyer and Seller subsidy.

In Col. (3), we examine heterogeneity based on personal ties. In 59% of pairs, both the buyer and seller said they had visited inside of the other's home prior to the start of the experiment. Even among such pairs (Panel B), the Seller subsidy leads to more irrigation than the Buyer subsidy (51% difference in treatment effects, p-value=0.096). In Col. (4), we examine heterogeneity by whether the buyer reported having an "excellent" or "very good" relationship with the seller before the start of the experiment. Among such pairs, which comprise 35.6% of the sample, we estimate that the Seller subsidy leads to substantially more trade than the Buyer subsidy (Panel B, 93% difference in effects, p-value=0.066).

We also see limited evidence that caste ties eliminate the contracting failure. The Seller subsidy group trades more than the Buyer subsidy group when both members of the pair are the same religion (Col. 5), same caste (Col. 6), or same subcaste (Col. 7).[29]

Overall, social or caste closeness do not appear sufficient to resolve the contracting friction—at least based on the proxies we have. However, sufficient social distance seems to undermine parties' ability to take advantage of the subsidy in the first place. For example, if the buyer and seller are different religions (i.e. a Hindu and a Muslim), the baseline level of trade among such pairs is low in the Control group, and such pairs do not make use of the subsidy offer (Col. 5, Panel A). Given the high returns to irrigation and the large size of the subsidy, this suggests that negative social ties may pose important barriers to trade in village communities.

5.4. **Future Trade.** In Table 9, we examine whether the one-shot subsidy had sustained impacts on trading behavior in the year after the experiment—after the subsidy was no longer in place. We see little evidence for sustained impacts: the amount of future trade among the Seller subsidy and Buyer subsidy groups is indistinguishable from that in the Control group, and from each other. This is true for both the full sample as well as the subset of pairs that had traded in the past.

The lack of effects on future trade is perhaps not surprising, given that treatment effects are present among both those that had traded in the past and those who have social ties. The results in Table 9 suggest that increased information on buyer types acquired during the experiment—for example, learning about buyers' trustworthiness after extending more credit to them in the Seller subsidy group—does not seem to be an operable margin, at least in terms of the impact of "subsidizing" experimentation among specific buyers and sellers through our intervention. More generally, these findings are in accordance with our interpretation that buyers and sellers choose not to enter into trade contracts due to perceived costs of future funds recovery.

5.5. **Division of Surplus.** By forgeoing trade in the Buyer-subsidy, pairs are giving up not only increased yields (for the buyer), but also the increase in irrigation sales (for the seller) and the subsidy payment itself. We can (roughly) estimate the difference in total surplus generated from trade across the treatments: Surplus = (Crop revenue) + (Subsidy payment) - (Cost to seller for irrigation) - (Crowd out of profits from sales with others). Applying this, we get a surplus difference between the Seller subsidy and Buyer subsidy of Rs. 5237.21. This corresponds to 16.1% of baseline annual household income among

[29]We asked each member of the pair to self-identify his religion, caste (organized in coarse categories for both Hindus and Muslims, such as upper caste, middle cast, OBC, SC, and ST), and his specific subcaste (based on a listing of subcastes in the state). In Col. (6), the covariate equals 1 if the parties have the same caste ranking, across Hindus and Muslims. The results are similar if we instead code the covariate as 1 only when both the religion and caste category are the same. In Col. (7), the covariate equals 1 if both the buyer and seller are in the exact same specific subcaste.

buyers.[30] As discussed above, this does not account for the cost of any potential complementary inputs which may have also increased to produce the yield increases. Consequently, we caution that this is best interpreted as an upper bound.

6. Survey Results — Funds Recovery in Village Factor Markets

Our simple intervention indicates that the anticipation of having to recover funds in the future from one's neighbor undermines mutually beneficial trade. This mechanism has the potential to affect trade in a broad range of economic activity in village economies. To supplement our experimental evidence, we conducted a short survey of sellers across four prominent agricultural factor markets: tractor rental, bullock rental, labor, and well rental for irrigation (i.e. the setting for our study). Each of these markets shares the features that we described for irrigation markets above: trade occurs through spot contracts in which farmers purchase factor inputs from fellow villagers, and deferred payments are infrequently observed (see below).

The goal of the survey was to ask sellers their beliefs related to contracting. Surveys were conducted among 124 sellers, randomly sampled from the population of each seller type across 13 villages (from a different set of villages than those in the experiment). The average respondent had been acting as a seller in his village for 13.1 years, and is therefore quite experienced.

We report survey results in Figure 7. The overwhelming majority of sellers (78%) have never had an agreement with a buyer before starting work, allowing him to defer payment (Panel A).

In Panel B, we ask respondents why sellers in their market do not generally enter into deferred payments, letting respondents agree with as many possible choices as they'd like as well as give reasons of their own. 90% of sellers state it is "difficult or costly to recover money later" as a contributing reason. Only 4% of respondents believe that sellers are okay with deferred payments.

In Panel C, sellers specify one source of costs: the perception that repeated attempts are needed to recover funds. When asked how many collection attempts would be needed to collect on a one-month deferred payment, only 21% of sellers believed that the buyer would deliver the payment himself as promised. In addition, only 2.4% of sellers believed that one trip by the seller would sufficient. Rather, 72% said repeated badgering would be needed, with about half the sellers predicting that it would take 3-10 visits to the buyer after the due date in order to get their money. These perceptions match the view summarized in Banerjee and Duflo (2007) about the village market for monetary credit: while the risk of absolute default seems low, this is achieved through high effort costs for funds recovery.

[30] To arrive at this estimate, we priced the cost to the seller for irrigation at its variable cost: the cost of the diesel that must be placed into the engine. If, instead, we price the cost of irrigation to the seller at the market price of irrigation (i.e. assuming the seller makes zero profits on any irrigation), these estimates fall slightly to 15.9% of annual income.

While village moneylenders appear to voluntarily enter this business, neighbors seem less likely to do so in our setting.

The question in Panel D suggests one reason why such recovery dynamics may have outsized costs for neighboring farmers. In addition to the time and effort required to visit a neighbor's home, sellers indicate that multiple recovery affect social relationships. 74% of respondents state that it is "likely" or "very likely" that multiple collection attempts strain the personal relationship between the buyer and seller. In our context, this may be heightened through the cultural feature that going to someone else's house and asking them for money (even if it is to recover money that is owed to oneself) is debasing.

The above results capture self-reported beliefs, and are therefore of course only suggestive. However, they indicate a set of views among sellers that accords with the pattern of revealed preference results we see in our experiment.

7. Conclusion

We study failures in intertemporal contracting among Indian farmers with neighboring landholdings. The Seller subsidy provides a benchmark for how much trade could occur in the presence of the subsidy. We find that the Buyer subsidy treatment falls far short of this benchmark. Moreover, the magnitude of effects on the amount of irrigation and yields indicate that this failure is economically meaningful for buyers. These findings suggest that, within the context of our experiment, contract enforceability is a first-order impediment to realizing the gains from trade. Given that so much economic activity in underdeveloped settings occurs through interpersonal exchange among co-villagers, our findings have potential relevance for understanding low investment and output in developing countries.

References

Anderson, Siwan. 2011. “Caste As an Impediment to Trade.” *American Economic Journal: Applied Economics.* 3(1): 239-263.

Banerjee, Abhijit and Esther Duflo. 2000. “Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry.” *Quarterly Journal of Economics.* 115(3): 989–1017.

Bardhan, Pranab K. 1983. “Labor-tying in a Poor Agrarian Economy: A Theoretical and Empirical Analysis.” *Quarterly Journal of Economics,* 501–514.

Belloni, Alenxandre, Victor Chernozhukov, and Christian Hansen. 2012. “Inference on Treatment Effects after Selection amongst High-Dimensional Controls.” *Review of Economic Studies*, 81(2): 608-650.

Belloni, Alenxandre, Victor Chernozhukov, and Christian Hansen. 2014. “High-Dimensional Methods and Inference on Structural and Treatment Effects.” *Journal of Economic Perspectives*, 28(2): 29-50.

Collins et al. 2010. *Portfolios of the Poor.* Princeton University Press.

Fisman, Raymond and Inessa Love. 2003. “Trade Credit, Financial Intermediary Development and Industry Growth.” *Journal of Finance*, 58(1): 353-374.

Gine, Xavier and Hanan Jacoby. 2015. “Markets, Contracts, and Uncertainty: A Structural Model of a Groundwater Economy.” Working Paper.

Iyer, Raj and Antoinette Schoar. 2008. “The Importance of Holdup in Contract Negotiations: Evidence from an Audit Study.” Working paper, MIT.

Jacoby, Hanan, Rinku Murgai and Saeed Rehman. 2004. “Monopoly Power and Distribution in Fragmented Markets: The Case of Groundwater.” *Review of Economic Studies.* 71: 783-808.

Machiavello, Rocco and Ameet Morjaria. 2014. “Competition and Relational Contracts: Evidence from Rwanda’s Coffee Mills.” Working Paper.

Machiavello, Rocco and Ameet Morjaria. 2016. “The Value of Relationships: Evidence from a Supply Shock to Kenya Rose Exports.” *American Economic Review*, forthcoming.

McMillan, John and Christopher Woodruff. 1999. “Interfirm Relationships and Informal Credit in Vietnam.” *Quarterly Journal of Economics.* 114(4): 1285-1320.

Meyerson, Roger and Mark Satterthwaite. 1983. “Efficient Mechanisms for Bilateral Trading.” Journal of Economic Theory. 29:265-281.

Townsend, Robert. 1994. “Risk and Insurance in Village India.” Econometrica, 62(3):539-591.

FIGURES

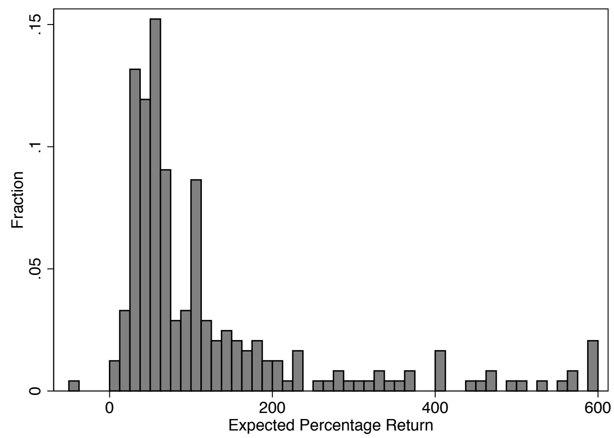

Figure 2. Perceived Returns to an Additional Irrigation
Notes: At baseline, water buyers were asked how much their revenues would increase if they did one more irrigation for each crop, and how much it would cost to purchase one more irrigation. The perceived percentage return is: [(revenue increase - irrigation cost)/irrigation cost]*100. The figure plots the average perceived return (averaged across all crops grown by the farmer) for each farmer. Values are topcoded at 600% in figure. N=243 water buyers in the experiment who were administered this question in the baseline survey.

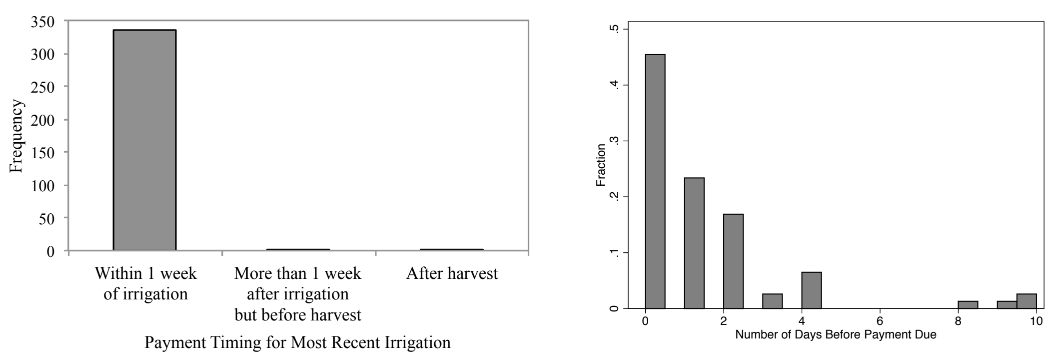

Figure 3. Payment Timing for Irrigation Purchases
Notes: The left panel plots water buyers' response to when they paid for their most recent irrigation purchase from each of their water sellers (N=337 buyer-seller pairs, from baseline survey). The right hand side panel plots data from the control group in the experiment sample, showing the number of days after irrigation when payment for irrigation was due within the experiment (N=77 transactions between buyer-seller pairs in the control group).

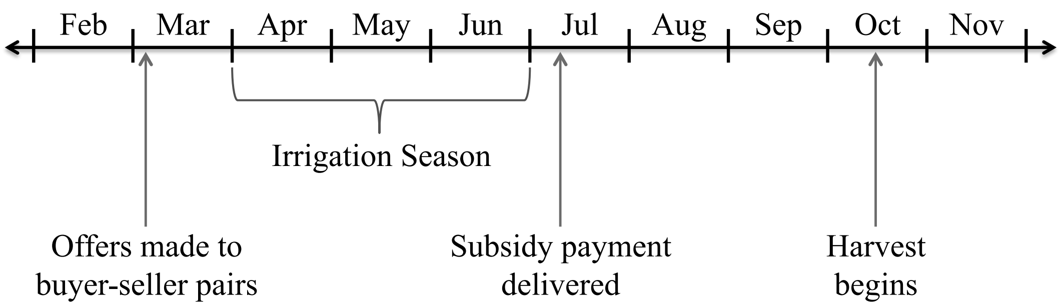

Figure 4. Experiment Timeline

Figure 5. Average Irrigation Levels by Week

Notes: This figure shows the amount of trade within buyer-seller pairs. It plots the raw average number of hours purchased in each week of the experiment, separately for each treatment group. The plot lines are smoothed, using a lowess smoother of 0.35.

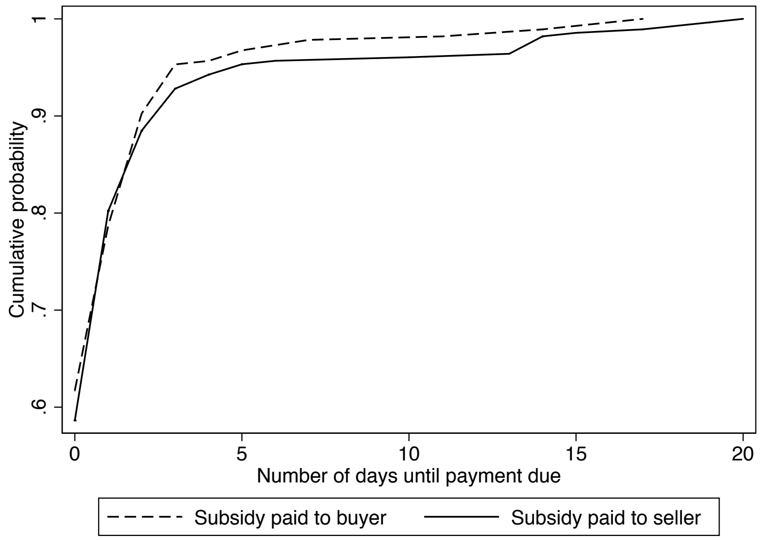

Figure 6. Trade Credit

Notes: The figure plots the CDF of number of days of trade credit for each of the two subsidy groups. Number of days before payment due = (Date when paymentwas made or expected) - (Date of irrigation), at time of weekly survey. Sample is restricted to pair-week observations where the pair irrigated.

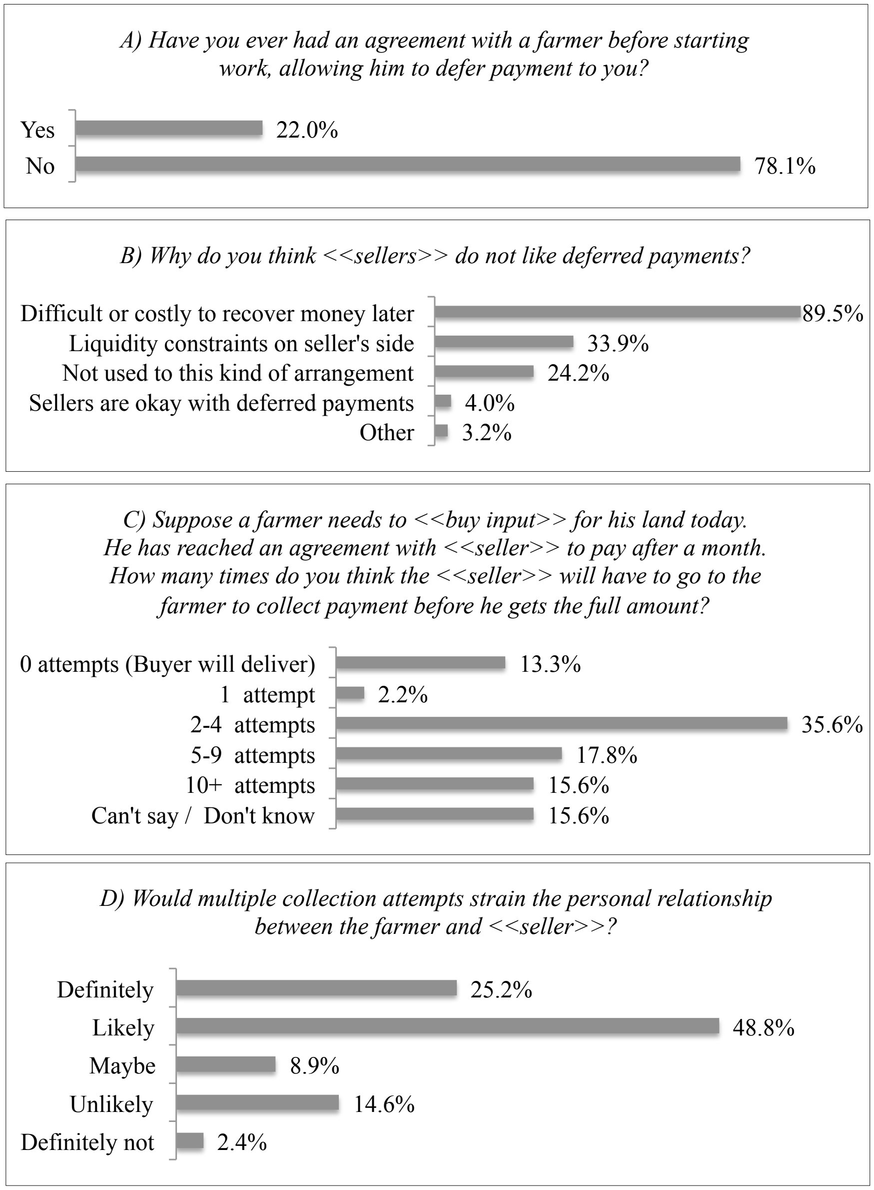

Figure 7. Intertemporal Exchange - Survey Responses

Notes: Survey conducted with sellers across four factor markets: bullock rental, tractor rental, well rental, and labor. N=124 sellers, chosen randomly from sellers in 13 villages in Odisha, India. The words in << >> were adapted to the factor sold by the respondent. E.g., in Panel B, the question was posed as "Why do you think <<plough owners / tractor owners / well owners / laborers>> do not like deferred payments?" in order to make questions concrete for respondents.

TABLES

Table 1. Summary Statistics and Balance

	Full sample		Seller subsidy		Buyer subsidy		Diff.
	Mean	Std dev	Coeff.	p-val.	Coeff.	p-val.	p-val
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Seller characteristics							
Value of consumption in past year (Rupees)	36772	61057	50	0.993	7155	0.535	0.484
Value of assets (Rupees)	336940	685106	-106561	0.405	-25203	0.871	0.349
Number of neighbors growing sugarcane (potential buyers)	4.068	1.500	0.113	0.614	-0.016	0.943	0.430
Number of neighboring plots with borewell	2.206	1.364	-0.196	0.282	-0.377	0.037	0.210
No other neighboring plots have borewell	0.189	0.392	0.045	0.345	0.074	0.123	0.497
Religion is Hindu	0.880	0.326	0.008	0.820	0.027	0.410	0.490
Seller is Scheduled caste/tribe	0.015	0.121	0.021	0.083	0.018	0.104	0.865
Buyer characteristics							
Value of consumption in past year (Rupees)	21128	8577	-3262	0.066	-2412	0.186	0.402
Value of assets (Rupees)	120183	106433	8249	0.652	-9564	0.447	0.243
Number of neighboring plots with borewell	2.987	1.415	0.022	0.905	0.017	0.926	0.976
No other neighboring plot has well (except assigned seller)	0.138	0.345	-0.003	0.947	0.023	0.596	0.487
Irrigated any plot in previous irrigation season	0.856	0.352	-0.039	0.520	-0.002	0.973	0.461
Spent money on irrigation at any point in past year	0.940	0.238	0.123	0.017	0.079	0.118	0.116
Religion is Hindu	0.850	0.357	0.022	0.597	0.035	0.387	0.692
Buyer is Scheduled caste/tribe	0.037	0.189	0.007	0.739	0.021	0.372	0.541
Social & market distance							
Buyer and seller have traded in past	0.624	0.485	0.047	0.451	0.026	0.674	0.683
Buyer and seller have both visited inside each other's home	0.590	0.492	0.158	0.017	0.136	0.040	0.693
Buyer and seller both report very good relationship	0.334	0.472	0.056	0.381	0.036	0.569	0.700
Buyer and seller are same religion	0.894	0.308	-0.018	0.667	-0.003	0.938	0.657
Buyer and seller are same caste category	0.661	0.474	-0.061	0.335	-0.057	0.352	0.949
Buyer and seller are same subcaste	0.486	0.500	-0.072	0.271	0.002	0.973	0.177
Buyer's consumption higher than seller's last year	0.265	0.442	-0.243	0.005	-0.138	0.118	0.092
Buyer has more assets than the seller	0.372	0.484	-0.042	0.605	-0.003	0.967	0.573

Notes: Cols (1)-(2) show sample mean and standard deviation for the full sample. In Cols. (3)-(6), each row reports results from a separate regression of the covariate on dummies for Seller subsidy and Buyer subsidytreatments (Assignment to control is omitted category), and fixed effects for each village (strata). P-values for each coefficient are based on robust standard errors. Col (7) reports the p-value of an F-test for whether the Seller subsidy treatment coefficient equals the Buyer subsidy treatment coefficient in each of these regressions. For covariates collected as part of the full baseline survey, conducted with a random 61% subsample of pairs, only those pairs are used for the estimates.

Table 2. Treatment Effects on Trade

	Dependent Variable: Hours of Irrigation				*Dependent Variable: Irrigated (dummy)*			
	Full sample (1)	Full sample (2)	Buyer & seller have traded in past (3)	Buyer & seller have not traded in past (4)	Full sample (5)	Full sample (6)	Buyer & seller have traded in past (7)	Buyer & seller have not traded in past (8)
Panel A: Irrigation Season								
Subsidy paid to water seller	0.799	1.471	2.461	0.081	0.0989	0.1118	0.1959	0.0183
	(0.310)***	(0.351)***	(0.495)***	(0.138)	(0.0235)***	(0.0227)***	(0.0273)***	(0.0649)
Subsidy paid to water buyer	0.270	0.822	1.539	0.102	0.0599	0.0624	0.1174	0.0105
	(0.279)	(0.330)***	(0.329)***	(0.122)	(0.0235)***	(0.0215)***	(0.0252)***	(0.0566)
Estimator	OLS	Tobit ME	Tobit ME	Tobit ME	OLS	Logit ME	Logit ME	Logit ME
P-value: Seller subsidy = Buyer subsidy	0.034	0.023	0.000	0.839	0.072	0.026	0.010	0.863
Observations (pair-weeks)	3,663	3,663	2,286	1,377	3,663	3,663	2,286	1,377
Dependent var mean: Buyer subsidy	1.174	1.174	1.564	0.528	0.150	0.150	0.199	0.069
Panel B: Full Subsidy Period								
Subsidy paid to water seller	0.578	1.190	1.971	-0.009	0.0701	0.0794	0.1379	0.0115
	(0.269)**	(0.325)***	(0.398)***	(0.101)	(0.0175)***	(0.0171)***	(0.021)***	(0.0476)
Subsidy paid to water buyer	0.189	0.640	1.193	0.076	0.0432	0.0438	0.0799	0.0108
	(0.242)	(0.308)***	(0.331)***	(0.108)	(0.0175)**	(0.0161)***	(0.0188)***	(0.0409)
Estimator	OLS	Tobit ME	Tobit ME	Tobit ME	OLS	Logit ME	Logit ME	Logit ME
P-value: Seller subsidy = Buyer subsidy	0.057	0.032	0.013	0.270	0.098	0.034	0.011	0.983
Observations (pair-weeks)	5,698	5,698	3,556	2,142	5,698	5,698	3,556	1,498
Dependent var mean: Buyer subsidy	0.947	0.947	1.298	0.367	0.116	0.116	0.156	0.050

Notes: The dependent variables are the number of hours of irrigation purchased by the buyer from his paired seller (Cols. 1-4) and a dummy for whether the buyer and seller traded (Cols. 5-8). Observations in Panel A are comprised of the irrigation season (May-June), and in Panel B all weeks when pairs were eligible to receive the subsidy (April-June). Cols. 3 and 7 (4 and 8) restrict analysis to pairs in which the buyer and seller had (not) traded irrigation with each other in the past. The omitted category in all regressions is Assignment to Control. All regressions contain village fixed effects and baseline covariate controls for crop type, wealth, caste, and market power. Each regression reports OLS estimates or Logit or Tobit marginal effects, as indicated. Standard errors clustered by buyer-seller pair.

Table 3. Crowd Out: Buyer's Transactions with Other Sellers

	Any other sellers (dummy)		Number of other sellers		Number of irrigations from other sellers		Irrigation hours			
							Other sellers	Paired seller	Total hours	Total hours
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Subsidy paid to water seller	-0.0131	-0.0302	-0.0159	-0.0163	-0.0255	-0.0470	-0.193	7.19	7.00	13.41
	(0.064)	(0.0642)	(0.077)	(0.0089)*	(0.345)	(0.059)	(0.392)	(2.94)**	(2.91)**	(3.72)***
Subsidy paid to water buyer	0.00127	-0.0116	0.00249	-0.0070	0.331	0.0478	0.169	2.43	2.60	6.94
	(0.062)	(0.064)	(0.074)	(0.0115)	(0.422)	(0.099)	(0.465)	(2.65)	(2.63)	(3.24)**
P-value of difference	0.779	0.695	0.774	0.338	0.392	0.066	0.294	0.044	0.062	0.055
Sample	Full sample	Full sample	Full sample	Full sample	Full sample	Full sample	Full sample	Full sample	Full sample	Previous trade
Estimator	OLS	Logit ME	OLS	Tobit ME	OLS	Tobit ME	OLS	OLS	OLS	OLS
Observations (pairs)	407	385	407	407	407	407	407	407	407	254
Dep var mean: Buyer subs	0.300	0.317	0.341	0.341	1.729	1.729	1.028	10.565	11.59	14.78

Notes: Cols. 1-6 examine measures of crowd-out: whether buyers change how much they trade with sellers other than their paired seller; these were captured directly from questions in the endline survey. In Col. 7, the dependent variable is computed as [(total payments made to a seller)/(average hourly price charged by that seller)], summed across all sellers except the buyer's paired seller. The dependent variable in Col. 8 is the total number of irrigation hours across the season within the pair (this aggregates the weekly treatment effects to the pair level). Cols. 9-10 estimate effects on total irrigation hours purchased from any seller (the paired seller + any other sellers). In Col. 10, the sample is limited to pairs that had traded prior to the experiment. The omitted category is Assignment to Control. Each regression reports OLS estimates or Logit or Tobit marginal effects, as indicated. All regressions contain village fixed effects and baseline covariate controls for crop type, wealth, caste, and market power. Robust standard errors.

Table 4. Treatment Effects on Yields

	Dependent variable: Yields index (std. dev)		*Dependent variable: Log crop revenue*	
	(1)	(2)	(3)	(4)
Subsidy paid to water seller	0.411	0.514	0.0968	0.214
	(0.202)**	(0.218)***	(0.0798)	(0.108)*
Subsidy paid to water buyer	0.076	0.098	0.00247	0.0553
	(0.193)	(0.218)	(0.0788)	(0.114)
P-val: Buyer subsidy = Seller subsidy	0.0467	0.0268	0.117	0.0598
Sample	Full sample	Previous trade	Full sample	Previous trade
Observations (pairs)	361	219	361	219

Notes: The dependent variable in Cols. 1-2 is a composite index of crop yields, normalized by the standard deviation of yields for each crop among the Control group. Crop revenue equals (self-reported yield levels * self-reported prices). The dependant variable in Cols. 3-4 log revenue, estimated as the inverse hyperbolic sine of crop revenue (due to zero revenue values for fallow plots). Cols. 2 and 4 restrict analysis to pairs in which the buyer and seller had traded with each other in the past. OLS regressions. All regressions include village fixed effects, the standard baseline controls, and crop dummies (sugarcane, wheat, rice, or fallow). Robust standard errors reported.

Table 5. Ex-Ante Transfers: Price Reductions

Dependent variable	Deviation from market price (1)	Price discount (dummy) (2)	Price discount (dummy) (3)
Subsidy paid to water seller	-15.22	0.0494	-0.0374
	(16.792)	(0.046)	(0.0965)
Subsidy paid to water buyer	4.585	-0.0197	-0.1381
	(14.026)	(0.047)	(0.1008)
P-value: Buyer subsidy = Seller subsidy	0.0777	0.0150	0.0120
Estimator	OLS	OLS	Logit ME
Observations (pair weeks)	627	627	412
Dependent var mean: Buyer subsidy group	4.29	0.084	0.123

Notes: Amount of price discount = (Market value - Amt charged). Market value = (modal price in village among control group)*hours of irrigation. Cols. (1)-(2) shows OLS estimates. Col. (3) reports estimated marginal effects from a logit regression. All regressions contain village fixed effects and baseline controls for crop type, wealth, caste, and market power. The sample is restricted to pair-weeks where the pair irrigated. There were 7 villages where no discounts were ever offered; the regression in Col. (3) drops observations where the village fixed effects and controls perfectly predict no discount. Standard errors are clustered by pair.

Table 6. Trade Credit

Dependent variable	Number of days before payment due (1)	Any deferred payment (2)	Any deferred payment (3)
Subsidy paid to water seller	0.0810	0.0213	0.0289
	(0.036)**	(0.009)**	(0.0101)***
Subsidy paid to water buyer	0.0135	0.00399	0.0060
	(0.032)	(0.009)	(0.00878)
P-value: Buyer subsidy = Seller subsidy	0.0456	0.0286	0.0080
Estimator	OLS	OLS	Logit ME
Observations (pair-weeks)	5,698	5,698	5,600
Dependent var mean: Buyer subsidy group	0.118	0.045	0.045

*Notes:*Number of days before payment due = (Date when payment was made or expected) - (Date of irrigation), at time of weekly survey. Each regression reports OLS estimates or Logit marginal effects, as indicated. All regressions contain village fixed effects and baseline controls for crop type, wealth, caste, and market power. Standard errors are corrected to allow clustering by pair.

Table 7. Heterogeneity — Wealth

Wealth proxy measure	*Dependent variable: Hours of irrigation*		*Dependent variable: Irrigated (dummy)*	
	Assets (1)	*Consump* (2)	*Assets* (3)	*Consump* (4)
	Panel A: Buyer has below median wealth			
Seller subsidy	0.899	1.298	0.0657	0.0938
	(0.505)***	(0.466)***	(0.0282)**	(0.0262)***
Buyer subsidy	-0.012	0.191	0.0120	0.0210
	(0.504)	(0.455)	(0.0260)	(0.0226)
P-value: Seller subsidy = Buyer subsidy	0.026	0.002	0.046	0.002
	Panel B: Buyer has above median wealth			
Seller subsidy	2.571	0.996	0.1576	0.0416
	(0.629)***	(1.053)	(0.0357)***	(0.0573)
Buyer subsidy	1.666	1.924	0.0847	0.1211
	(0.604)**	(0.778)**	(0.0311)***	(0.0497)**
P-value: Seller subsidy = Buyer subsidy	0.141	0.387	0.102	0.258

Notes: This table shows treatment effects separately by the buyer's baseline wealth level. Two proxies for wealth are used: the total value of all household assets (Cols. 1 and 3), and the total value of consumption in the past year (Cols. 2 and 4). The dependent variable in Cols. 1-2 is the number of hours of irrigation and in Cols. 3-4 is a binary indicator for whether the pair irrigated that week. Panel A (B) reports estimated treatment effects of each treatment, relative to Control, for the case where the buyer has a below (above) median value of the indicated wealth proxy. In each column, the displayed coefficients are estimated marginal effects (computed using Stata's margins command) from a tobit regression on the full sample in Cols. (1)-(3) and a logit regression on the full sample in Cols. (4)-(6). Regressions include village fixed effects and baseline covariate controls for crop type, wealth, caste, and market power. Standard errors are clustered by pair. N=5,698 pair-weeks.

TABLE 8. Heterogeneity — Potential Correlates of Informal Enforcement Mechanisms

Covariate	*Market power*		*Social distance*						
	Buyer has no other potential sellers (1)	*Seller has no other potential buyers* (2)	*Buyer & seller visited each other's home* (3)	*Buyer & seller have very good relationship* (4)	*Buyer & seller are same religion* (5)	*Buyer & seller have same caste ranking* (6)	*Buyer & seller are the same subcaste* (7)	*Seller is lower caste than buyer* (8)	*Buyer has higher consumption than seller* (9)
Covariate mean (all observations)	0.138	0.189	0.590	0.356	0.894	0.661	0.486	0.162	0.265
Panel A: Covariate=0									
Seller subsidy	1.221	1.253	0.300	0.860	-0.320	1.153	1.490	1.271	1.190
	(0.339)***	(0.352)***	(0.392)	(0.385)**	(1.037)	(0.607)*	(0.470)***	(0.342)***	(0.581)**
Buyer subsidy	0.530	0.446	-0.319	0.496	-0.360	0.598	1.100	0.661	0.139
	(0.335)	(0.349)	(0.410)	(0.371)	(0.892)	(0.576)	(0.482)**	(0.341)*	(0.619)
P-val: Seller subs = Buyer subs	0.012	0.005	0.103	0.217	0.961	0.206	0.263	0.027	0.013
Dep var mean: Control group	0.676	0.649	0.562	0.760	0.366	0.411	0.271	0.768	0.983
Panel B: Covariate=1									
Seller subsidy	0.798	0.810	1.620	1.846	1.360	1.211	1.046	0.642	2.228
	(0.866)	(0.670)	(0.489)***	(0.594)***	(0.344)***	(0.389)***	(0.470)**	(1.043)	(0.811)***
Buyer subsidy	1.212	1.343	1.075	0.957	0.763	0.663	0.342	0.610	1.573
	(0.793)	(0.645)**	(0.477)**	(0.536)*	(0.325)**	(0.363)*	(0.417)	(0.968)	(0.633)**
P-val: Seller subs = Buyer subs	0.560	0.362	0.092	0.066	0.029	0.075	0.049	0.963	0.433
Dep var mean: Control group	0.877	1.018	0.867	0.580	0.739	0.824	1.114	0.236	0.718

Notes: The dependent variable is the number of irrigation hours purchased by the buyer from his paired seller. Each column shows treatment effects separately for the binary baseline covariate that is listed at the top of the column. The overall sample mean of the covariate is reported at the top of the table. Panel A (B) reports estimated treatment effects of each treatment, relative to Control, for the case where the covariate equals 0 (equals 1). The bottom od Panel A (B) reports the dependent variable mean among Control group pairs when the covariate equals 0 (equals 1). The covariate in Col. (3) is a dummy that equals 1 if the buyer and seller each report having visited inside the other's home prior to the start of the experiment. The covariate in Col. (4) is a dummy that equals 1 if the buyer reported having an excellent or very good relationship with the seller before the experiment. The covariate in Col. (9) is a dummy that equals 1 if the buyer had higher annual consumption than the seller last year. In each column, the displayed coefficients are estimated marginal effects (computed using Stata's margins command) from a tobit regression on the full sample. Regressions include village fixed effects and baseline covariate controls for crop type, wealth, caste, and market power. Standard errors clustered by pair. N=5,698 pair-weeks.

Table 9. Trade in the Following Year

	Dependent variable: Hours of Irrigation				Dependent variable: Number of Irrigations			
Sample	Full sample	Full sample	Previous trade	Previous trade	Full sample	Full sample	Previous trade	Previous trade
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Subsidy paid to water seller	2.010	1.362	3.233	0.998	0.248	0.207	-0.165	0.036
	(3.60)	(2.431)	(4.814)	(5.382)	(0.328)	(0.260)	(0.480)	(0.558)
Subsidy paid to water buyer	-1.402	-1.364	0.054	-0.903	-0.143	-0.126	-0.566	-0.216
	(3.52)	(2.378)	(4.900)	(3.346)	(0.326)	(0.258)	(0.471)	(0.308)
P-value: Buyer subsidy = Seller subsidy	0.200	0.133	0.409	0.927	0.128	0.098	0.240	0.948
Estimator	OLS	Tobit ME	OLS	Tobit ME	OLS	Tobit ME	OLS	Tobit ME
Observations (pairs)	407	407	254	254	407	407	254	254
Dependent var mean: Buyer subsidy	19.35	19.35	20.74	20.74	2.51	2.51	2.51	2.51

Notes: The dependent variables capture the total amount of trade between the buyer and his paired seller during the irrigation year after the experiment ended. All regressions include village fixed effects and baseline covariate controls for crop type, wealth, caste, and market power. Robust standard errors.

Appendix A. Supplemental Figures and Tables

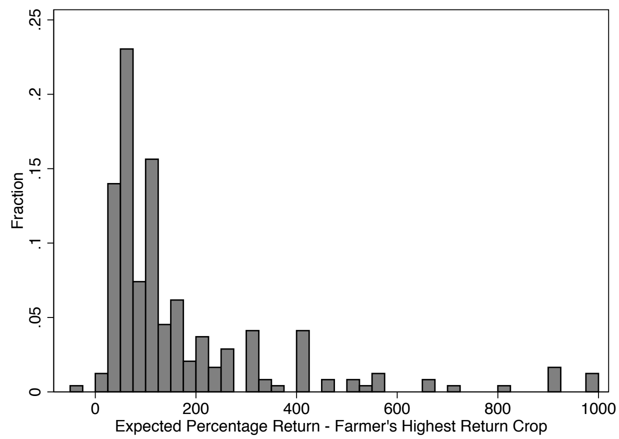

Figure A1. Perceived Returns to an Additional Irrigation

Notes: At baseline, water buyers were asked how much their revenues would increase if they did one more irrigation for each crop, and how much it would cost to purchase one more irrigation. The perceived percentage return is: [(revenue increase - irrigation cost)/irrigation cost]*100. The figure plots the perceived return for the crop identified by each respondent as having the highest return to irrigation. Values are topcoded at 1000% in figure. N=243 water buyers in the experiment who were administered this question in the baseline survey.

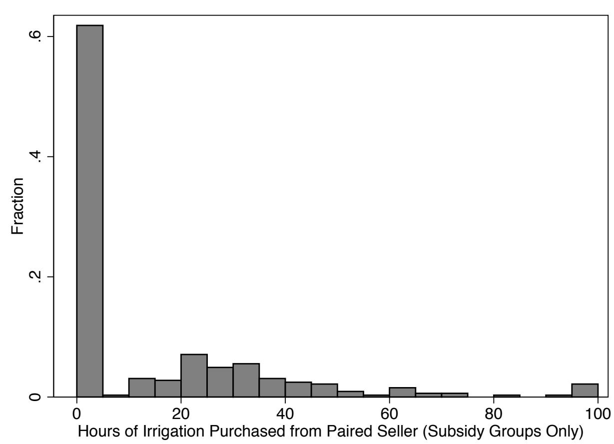

Figure A2. Distribution of Irrigation Hours

Notes: This figure shows the distribution of the total number of hours of irrigation purchased within buyer-seller pairs during the experimental period. The number of hours is topcoded at 100.

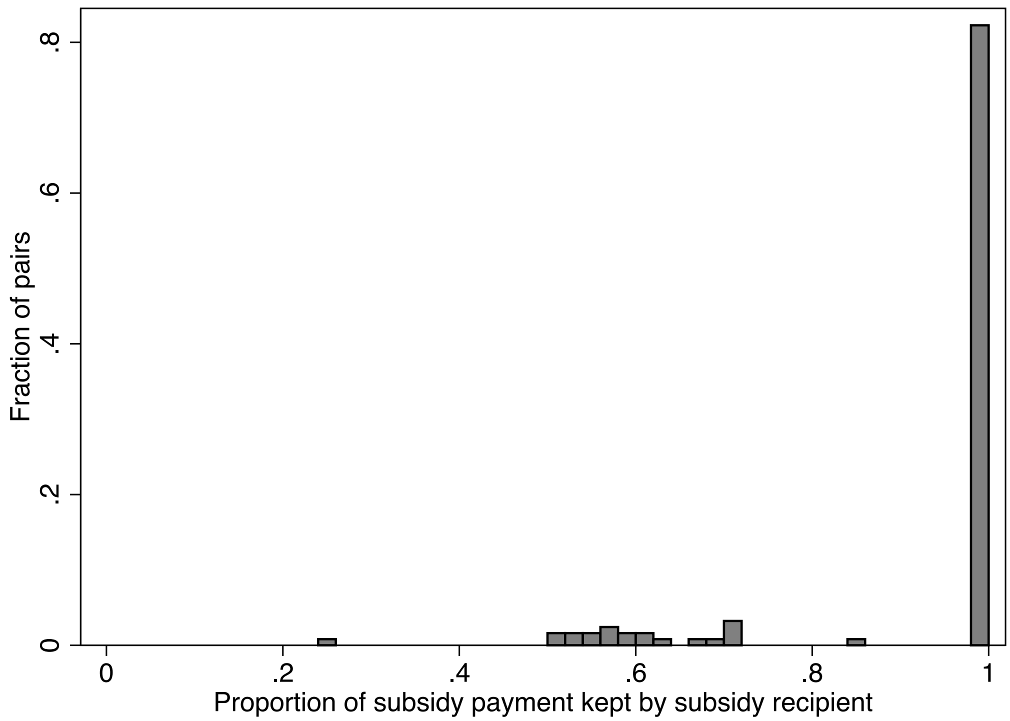

FIGURE A3. Ex-post Transfers: Sharing of the Subsidy

Notes: The sample is limited to pairs in the subsidy treatments who irrigated at least once and earned a subsidy (124 pairs). When the subsidy was delivered, each member of the pair was asked how the subsidy payment would be shared among them. The figure plots the distribution of the proportion of the subsidy payment that would be kept by the subsidy recipient (i.e. the seller in the Seller subsidy pairs and the buyer in the Buyer subsidy pairs).

Table A1. Robustness Checks

Dependent variable: Hours of Irrigation within Pair

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Subsidy paid to water seller	0.731**	0.736**	0.799**	0.784**	0.799**	0.799**	0.805**
	(0.314)	(0.314)	(0.310)	(0.288)	(0.310)	(0.317)	(0.312)
Subsidy paid to water buyer	0.352	0.275	0.270	0.323	0.270	0.270	0.298
	(0.273)	(0.284)	(0.279)	(0.272)	(0.279)	(0.285)	(0.285)
All covariates in balance table?	No	Yes	No	No	No	No	No
Crop, wealth, caste, & market power controls?	No	No	Yes	No	Yes	Yes	Yes
Post-LASSO controls?	No	No	No	Yes	No	No	No
Calendar week fixed effects?	No	No	No	No	Yes	No	No
Village x Calendar week fixed effects?	No	No	No	No	No	Yes	No
Reported yields?	No	No	No	No	No	No	Yes

Notes: OLS regressions. Sample is pair-weeks during the irrigation season. All regressions have fixed effects for each village (strata). Col. (2) includes all covariates reported in the balance table. Col. (3) includes baseline controls for crops grown, along with the three primary sources of heterogeneity examined in the paper: wealth, caste distance, and market power. Col. (4) includes controls selected using the post-double selection method, using LASSO to select controls (note controls selection did not include a clustering correction, on which econometric guidance was not available). Col. (5) adds week fixed effects and Col. (6) adds village*week fixed effects to the basic crop and heterogeneity controls. Standard errors clustered by pair.

Table A2. Crowd-Out: Seller's Transactions With Other Buyers

	Any other buyers (dummy)		Number of other buyers		Irrigations to other buyers	Revenue from other buyers	Irrigation hours			
							Other buyers	Paired buyer	Total	Total
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Subsidy paid to water seller	-0.0791	-0.0705	-0.119	-0.0874	-0.271	-18.55	-0.536	7.19	6.65	11.61
	(0.054)	(0.063)	(0.067)*	(0.0132)***	(0.319)	(97.515)	(1.677)	(2.94)**	(3.11)**	(3.86)**
Subsidy paid to water buyer	0.000328	0.0198	-0.000650	0.0179	0.394	40.92	0.699	2.43	3.13	5.13
	(0.055)	(0.0650)	(0.074)	(0.0187)	(0.418)	(99.088)	(1.732)	(2.65)	(2.90)	(3.36)
P-value of difference	0.0604	0.052	0.0250	0.000	0.0259	0.459	0.372	0.0444	0.186	0.0530
Sample	Full sample	Full sample	Full sample	Full sample	Full sample	Full sample	Full sample	Full sample	Full sample	Previous trade
Estimator	OLS	Logit ME	OLS	Tobit ME	OLS	OLS	OLS	OLS	OLS	OLS
Observations (pairs)	407	334	407	407	407	407	407	407	407	254
Dep var mean: Buyer subsidy	0.194	0.237	0.241	0.241	1.25	254.18	4.32	10.56	14.88	14.95

Notes: Cols. 1-6 examine measures of crowd-out: whether sellers change how much they trade with buyers other than their paired buyer; these were captured directly from questions in the endline survey. In Col. 7, the dependent variable is computed as [(total payments received from a buyer)/(average hourly price charged to that buyer)], summed across all buyers except the seller's paired buyer. The dependent variable in Col. 8 is the total number of irrigation hours across the season within the pair (this aggregates the weekly treatment effects to the pair level). Cols. 9-10 estimate effects on total irrigation hours sold to any buyer (the paired buyer + any other buyers). In Col. 10, the sample is limited to pairs that had traded prior to the experiment. The omitted category is Assignment to Control. Each regression reports OLS estimates or Logit or Tobit marginal effects, as indicated. All regressions contain village fixed effects and baseline covariate controls for crop type, wealth, caste, and market power. Robust standard errors.

Table A3. Correlation Between Treatments and Crop Choice

	Dependent variable: Grew sugarcane (dummy)			
	(1)	(2)	(3)	(4)
Subsidy paid to water seller	-0.0334	-0.0322	0.0301	0.0310
	(0.070)	(0.066)	(0.0977)	(0.0972)
Subsidy paid to water buyer	-0.0840	-0.0922	-0.007249	0.000428
	(0.070)	(0.067)	(0.0883)	(0.0879)
P-val: Buyer subsidy = Seller subsidy	0.359	0.260	0.607	0.660
Sample	Full sample	Full sample	Previous trade	Previous trade
Estimator	OLS	Logit ME	OLS	Logit ME
Observations (pairs)	361	347	219	199

Notes: The dependant variable is an indicator for whether the buyer grew sugarcane (the main cash crop). Cols. 3-4 restrict analysis to pairs in which the buyer and seller had traded with each other in the past. Each regression reports OLS estimates or Logit marginal effects, as indicated. All regressions include village fixed effects and the standard baseline covariate controls. Robust standard errors reported.

Table A4. Non-collection of Yields

	Coefficient (1)	Std error (2)
Panel A: Seller characteristics		
Value of consumption in past year (Rupees)	-7.48e-08	(0.000)
Value of assets (Rupees)	-3.69e-09	(0.000)
Number of neighbors growing sugarcane (potential buyers)	-0.00372	(0.012)
Number of neighboring plots with borewell	-0.0133	(0.012)
No other neighboring plots have borewell	-0.0131	(0.043)
Religion is Hindu	0.0220	(0.048)
Panel B: Buyer characteristics		
Value of consumption in past year (Rupees)	-7.05e-07	(0.000)
Value of assets (Rupees)	3.43e-07	(0.000)
Number of neighboring plots with borewell	-0.0165	(0.012)
No other neighboring plot has well (except assigned seller)	0.0542	(0.058)
Irrigated any plot in previous irrigation season	0.0290	(0.045)
Spent money on irrigation at any point in past year	-0.0371	(0.093)
Religion is Hindu	-0.00852	(0.047)
Panel C: Social & market distance		
Buyer and seller have traded in past	0.0511	(0.031)
Buyer and seller have both visited inside each other's home	-0.00972	(0.033)
Buyer and seller both report very good relationship	0.0653	(0.036)
Buyer and seller are same religion	0.0717	(0.040)
Buyer and seller are same caste category	0.0215	(0.032)
Buyer and seller are same subcaste	0.0386	(0.031)
Buyer's consumption higher than seller's last year	-0.102	(0.046)*
Buyer has more assets than the seller	0.0583	(0.056)
Panel D: Treatment indicators		
Subsidy paid to seller	-0.0118	(0.047)
Subsidy paid to buyer	-0.0601	(0.045)

Notes: Each row shows the results of a regression in which the dependent variable is an indicator for whether the buyer has a missing yields observation. In Panels A-C, the indicator is regressed on the covariate shown in that row. In Panel D, it is regressed on dummies for the Seller subsidy and Buyer subsidy treatments together (so that the ommitted category is the Control group). All regressions include village fixed effects. Co1. (1) reports the coefficient and Col. (2) reports the standard error from each regression. Robust standard errors. A coefficient of the form 7.48e-08 means 7.48*10^-8.