

Using individual-level randomized treatment to learn about market structure*

Lorenzo Casaburi Tristan Reed

March 2021

Abstract

Interference across competing firms in field experiments can be informative about market structure. An experiment that subsidizes a random subset of traders who buy cocoa from farmers in Sierra Leone illustrates this idea. Interpreting treatment-control differences in prices and quantities purchased from farmers through a standard model of Cournot competition reveals differentiation between traders is low. Combining this result with quasi-experimental variation in world prices shows that the number of traders effectively competing is 40% higher than the number operating in a village, highlighting the challenge of defining market boundaries. In turn, own-price and cross-price supply elasticities are high. These estimates suggest farmers face a competitive market.

JEL Classification: O12, O13, Q13, L13

Keywords: Field experiments, agricultural markets, market structure, differentiation, interlinked transactions.

*Previous versions of this paper were circulated with the titles “Interlinked Transactions and Pass-Through: Experimental Evidence from Sierra Leone” and “Competition and Interlinkages in Agricultural Markets: An Experimental Approach.” We thank Philippe Aghion, Pol Antràs, David Atkin, Dave Donaldson, Pascaline Dupas, Fred Finan, Matthew Gentzkow, Robert Gibbons, Rachel Glennerster, Oliver Hart, Asim Khwaja, Michael Kremer, Rocco Macchiavello, Ted Miguel, Ben Olken, Dina Pomeranz, Ori Shelef, Tavneet Suri, Chris Udry, Eric Verhoogen, Jack Willis, Josef Zweimüller and workshop participants at CEPR/LSE/TCD Development Economics Workshop, CSAE Oxford, the Edinburgh Conference on Agriculture and Structural Transformation, Harvard/MIT, LSE/UCL, the Montreal Workshop on Productivity, Entrepreneurship and Development, NBER Development Meeting, NBER Development and Organizational Economics Workshop, Paris School of Economics, Stanford, Stockholm University, Trinity College Dublin, UC Berkeley, UC San Diego, University of Naples, and University of Zurich for helpful suggestions and comments. Derick Bowen, Grant Bridgman, Felix Kanu and Fatoma Momoh provided excellent research assistance. We gratefully acknowledge the financial support of the International Growth Centre and the Agricultural Technology Adoption Initiative, and the institutional support of Innovations for Poverty Action in Freetown. Email: lorenzo.casaburi@econ.uzh.ch, treed@worldbank.org.

1 Introduction

A common challenge in experimental research is interference across treatment units, when the treatment status of an individual affects the outcomes of others. For example, a treatment that subsidizes some firms may change prices and quantities of other firms in the same market. In this case, the average differences between treatment and control firms in these outcomes cannot be interpreted as average treatment effects. To deal with this issue, researchers have typically employed *market-level* randomized treatments, assuming that the stable unit treatment value assumption (SUTVA) holds across markets rather than individuals. This approach however requires one to delineate market boundaries that are sufficiently large to incorporate all competitive spillovers and that enough markets are available to achieve statistical power (Muralidharan and Niehaus, 2017). In practice it is often infeasible for logistic or budgetary reasons.

In this paper, we show that differences in treatment and control prices stemming from *individual-level* randomized treatment can nonetheless be used to estimate key parameters of the industry equilibrium, namely differentiation among firms, supply (or demand) slope, and own-price and cross-price elasticities. The parameters are crucial for measuring market power and analyzing counterfactual equilibria, as in the industrial organization literature (see, e.g., Berry, Gaynor and Scott Morton, 2019).

We apply this approach to study competition among buyers in an agricultural export value chain. In the field experiment, we randomly subsidize a subset of traders buying raw cocoa (cacao) from farmers in Sierra Leone. The experiment induces variation in the traders' marginal revenue: a random sample of traders are paid a per-unit bonus for delivering high quality cocoa to five major wholesalers operating in the cocoa producing region of Sierra Leone. The bonus is equal to approximately 5% of wholesale price that traders receive.

We interpret the experiment through the lens of a standard model of Cournot competition among traders. Subsidizing a random subset of traders in the same market affects prices paid to farmers and quantities purchased by both treatment and control traders. Closed-form solutions show that the average difference in prices paid by treatment and control traders recovers the degree of differentiation among traders, which summarizes a trader's ability to buy while paying farmers a lower price than competitors.¹ Differentiation of buyers from the perspective of the seller could stem from a variety of factors such as reliability of demand, timeliness of payment, the buyers' ability to provide credit, search costs, or other idiosyncratic aspects of the relationship such as co-ethnicity,

¹Formally, the differentiation is defined as one minus the ratio between the slopes of a trader's inverse supply to own quantity and to each competitor's quantity.

friendship or trust. Intuitively, when traders are undifferentiated from the farmer’s perspective, the law of one price must hold. Additionally, the treatment-control difference in quantities identifies the slope of the trader inverse supply curve.

In the experiment, treatment and control traders pay similar prices. However, subsidized traders are eighty percent more likely to provide advance payments (27% vs 15%), suggesting that estimates of trader differentiation may vary once accounting for non-price competition. To account for it, we combine these results into an average treatment and control difference in the *effective price*, which describes the net present value of all payments to the farmer. To compute the effective price we follow several strategies, relying on cross-sectional variation in price and advances (our baseline strategy), heterogeneity in treatment responses along these two margins, rural banks’ interest rates, and an auxiliary experiment to measure farmers’ time preferences. In our baseline specification, treated traders pay an effective price higher than control traders by about 8.5% of the subsidy value, and results are qualitatively similar across different methods. In addition, treatment traders, who represent about 20% of the traders in the markets, purchase more than four times as much cocoa as control traders.

Due to competitive spillovers, one cannot interpret these average treatment and control differences as average treatment effects. Nevertheless, randomization is still essential to interpret the results through the lens of the model, as it ensures subsidies are uncorrelated with trader characteristics and thus that the treatment-control price differences arise only because of the subsidy and not underlying differences in cost. Our preferred estimate suggests that overall traders are relatively homogeneous, with a low value of the differentiation parameter, 0.09 on a 0-1 scale with 90% confidence interval of [-.11,.47]. The slope of the inverse supply is 0.35 [.24,.48]. These results are quite robust to alternative approaches to compute the effective price, with the highest differentiation rate (0.22) found when using data from the auxiliary experiment that measures farmers’ required rate of return.

By combining these parameters with the number of firms competing for farmers’ supply, one can compute own-price and cross-price supply elasticities and undertake counterfactual analysis. However, the number of firms in the market may not be directly observed. In our setting, traders are highly mobile across space and so is not clear how many bid for a given farmer’s output. Nonetheless, it is possible to infer the effective number of firms in the market by combining the experimental results with an estimate of the pass-through rate to farmers of a change in the wholesale price affecting all traders in the market, not just those receiving the experimental treatment. The world price is used as an instrument for the wholesale price, yielding an estimate of the pass-through

that is quite high, or 0.91, in line with other studies of the regional cocoa industry, e.g., Gayi and Tsowou (2015).

The number of competitors for a given farmer’s supply implied by this pass-through rate and the trader differentiation rate is 11.1 [8.8,18.9]. This value is 40% higher and significantly different (applying a 90% confidence interval) from the average number of traders we count operating in a village. This result suggests that markets cannot easily be delineated by village boundaries because farmers can sell their cocoa to traders outside their village. In supplementary analysis, we rule out alternative models of conduct (e.g., monopsony pricing) and we show that alternative moments from the data give similar values of differentiation and of the number of firms, providing additional support for the model.

Relatively low differentiation, high pass-through of common shocks, and large number of effective competitors suggest this market is highly competitive. Analysis of supply elasticities derived from these parameters supports this claim. The own-price supply elasticity, or the percent increase in quantity a trader experiences in response to a one percent increase in own price, holding constant the prices of all the other competitors, is very high, at 324 (in the limit case of perfect competition, with atomistic firms, this elasticity would be infinite). The cross-price supply elasticity is also large, at -32: if one of the competitors increases the price by 1% and competitors do not respond, each of the other 10 competitors loses about one third of its supply. The competitiveness of the market is also illustrated by an estimate of the impact of the experimental subsidy on aggregate prices and quantities, relative to a counterfactual without the experiment. The experimental subsidy induced an increase in price of 25% of the subsidy value for treatment traders and of 16.5% for control ones. Treatment traders increased supply by 174% and control traders reduced it by 38%, suggesting that the majority (90%) of the treatment traders’ increase in quantity comes from stealing control traders’ market share, rather than from an increase in aggregate supply.

The paper contributes to a literature on the implications of interference in experiments, sometimes called ‘spillovers’ or ‘externalities’, which occurs when the treatment status of an individual affects the outcomes of others (Baird, Bohren, McIntosh and Özler, 2018). Several recent experimental and quasi-experimental studies use market-level variation in individual-level treatment to study competition (Busso and Galiani, 2019, Jensen and Miller, 2018, Mitra, Mookherjee, Torero and Visaria, 2018, Rotemberg, 2019. Bergquist and Dinerstein, 2020).² An implicit assumption

²Other recent experiments (or natural experiments) featuring market-level variation include Crépon, Duflo, Gurgand, Rathelot and Zamora (2013), Cunha, De Giorgi and Jayachandran (2018), Lalive, Landais and Zweimüller (2015), Hildebrandt, Nyarko, Romagnoli and Soldani (2015), Burke, Bergquist and Miguel (2018), Filmer, Friedman, Kandpal and Onishi (2018), Mobarak and Rosenzweig (2014), Muralidharan, Niehaus and Sukhtankar (2016), Breza and Kinnan (2016). In non-experimental studies of competition, industry equilibrium models are typically estimated

in this strand of work is that the SUTVA required for standard experimental analysis holds when comparing markets. Donaldson (2015) argues this assumption is unlikely to hold if markets are integrated through trade. A potential response is to work with very large randomization clusters that encompass most of the spillovers (see, e.g., Muralidharan and Niehaus, 2017; Burke, Bergquist and Miguel, 2018). However, such an approach is often either very expensive or not logistically feasible because the number of such clusters is too small. Our alternative approach relies on violations of the SUTVA and individual-level randomized variation in firms’ marginal revenue to estimate the parameters of a standard industry equilibrium model. While our focus is a model of supply to potentially differentiated buyers, similar arguments combined with experimental variation in costs could be used to estimate a model of demand for potentially differentiated products (see e.g., Vives, 2001).

Our approach has broad application in policy analysis. Weyl and Fabinger (2013) highlights the pass-through rate as an economic tool for assessing the incidence of taxes or subsidies affecting all firms in a market. An estimate of the differentiation rate allows one to study further the incidence of taxes or subsidies given to only a subset of firms. Many industrial and agricultural policies are implemented in this way. For example, while in West Africa governments have broadly preferred universal subsidies for fertilizer use, in East Africa subsidies are only available to poorer farmers (Druilhe and Barreiro-Hurle, 2012). Outside of agriculture, subsidies often target only small or medium-sized firms (Chatzouz, Gereben, Lang and Torfs, 2017; Rotemberg, 2019), exporters (Rodrik, 1993; Panagariya, 2000; Kalouptsidi, 2018), or the politically connected (Khwaja and Mian, 2005; Faccio, 2006; Rijkers, Baghdadi and Raballand, 2015). Intuitively, as we demonstrate through counterfactual analysis, when subsidies are offered only to a subset of firms, the equilibrium effect of the subsidy depends both on the number of firms in the market and the extent to which they are differentiated. Identification of these parameters requires both an estimate of the market-level pass-through rate and of the average difference in prices paid by treatment and control firms.

Our estimate of buyer differentiation informs a literature on credit transactions in the context of agriculture, which are described as interlinked when credit is provided by buyers and output prices are co-determined with the interest rate (see, e.g., Bardhan, 1980; Bell, 1988).³ We find

using a time series of prices (and costs) and instrumental variables (Graddy, 1995; Osborne, 2005).

³Casaburi and Willis (2018), Casaburi and Macchiavello (2019), Ghani and Reed (2021), and Macchiavello and Morjaria (2020) provide recent empirical contributions on interlinked credit transactions in agricultural value chains in sub-Saharan Africa. Emran et al. (2020) documents the importance of financing middlemen in Bangladesh edible oils supply chain. More broadly, interlinked credit is a feature of all economies in the form of trade credit, which is an important source of finance for smaller firms (Petersen and Rajan, 1997), certain industries (Fisman and Love, 2003), and for international trade (Antras and Foley, 2015). In the United States, non-farm enterprises with fewer than 500 employees rely on trade credit for about 60% of their external finance (Mach et al., 2006).

that, in response to subsidies, treated traders use credit to farmers to secure supply. The advance payment is repaid when the farmer accepts a below market price for output at harvest time. While other studies have emphasized that credit supply is diminished in the absence of the market power required to sustain a relational credit contract (McMillan and Woodruff, 1999; Macchiavello and Morjaria, 2020), we demonstrate that credit provision need not coincide with high market power of credit providers. There are two plausible explanations for this finding. First, credit contract enforcement through customary legal institutions (i.e., the Paramount Chieftaincy described by Acemoglu, Reed and Robinson (2014)) may be effective, obviating the need for relational credit contracts. Second, as in the analysis of African trade credit relationship by Fisman and Raturi (2004), competition may increase borrowers’ incentives to establish creditworthiness, by reducing holdup concerns on the lender side.

Finally, our results contribute to a substantial literature describing the industrial organization of agricultural markets in low- and middle-income countries. Prior studies of farm-gate buyers’ market power have relied on primarily observational, rather than experimental evidence, analyzing trader price-cost margins (for Sub-Saharan Africa, see, e.g., Fafchamps, Gabre-Madhin and Minten, 2005; Osborne, 2005; Sitko and Jayne, 2014), price dispersion across space (Fackler and Goodwin, 2001; Aker, 2010), or the pass-through of international prices along the supply chain (Fafchamps and Hill, 2008; Dillon and Barrett, 2015).⁴ To the best of our knowledge, this is the first experiment that randomized any treatment at the trader level in agricultural markets and that used experimental subsidies to study competition. Our finding that farm-gate buyers in Africa have more limited market power consistent with evidence in the recent review by Dillon and Dambro (2017). Bergquist and Dinerstein (2020) and Iacovone and McKenzie (2019) study experimentally how vendors (of maize and fresh produce, respectively) adjust consumer prices in response to market-level subsidies affecting marginal costs. In this retail setting, these authors find a high degree of market power.

2 A Simple Model of Strategic Competition between Treatment and Control Firms

This section presents the model of imperfect competition we use to interpret our experimental results. The model illustrates how a subsidy (treatment) to a randomly-selected subset of firms in the market affects prices and quantities of all firms, not just the treated ones. In other words, the

⁴Chatterjee (2020) uses price dispersion across space to study crop markets in India. More broadly, recent theoretical and empirical contributions on the role of intermediaries in supply chains include Antràs and Costinot (2011), Bardhan et al. (2013), Chau et al. (2016), Maitra et al. (2017), and Emran et al. (2020).

SUTVA fails. Given our empirical setting, we focus on demand-side (i.e., trader) oligopsonistic competition, though a similar approach could be used to estimate a model with supply-side oligopolistic competition. The model is stylized and makes important assumptions, e.g. about conduct and functional forms. Since we use our empirical results to validate these assumptions, we postpone to Section 5 a detailed discussion of these assumptions and their implications.

The model provides a transparent closed-form mapping between individual-level treatments and three parameters describing the industry equilibrium. First, the average difference in prices paid to farmers by treatment and control traders identifies the degree of differentiation among traders. This parameter is defined as one minus the ratio between the slopes of a trader’s inverse supply to own quantity and to each competitor’s quantity. It thus measures the extent to which each trader is “insulated” from its competitors. Intuitively, without differentiation, the law of one price must hold and there is no difference in price paid by treatment and control traders.⁵

Second, for given differentiation, the average difference in quantities between treatment and control traders identifies the slope of the inverse supply curve facing the trader.

Third, for given differentiation and slope of inverse supply, the market-level pass-through rate (i.e., the response of the farmer price to a change in the wholesale price common to all traders) identifies the number of firms effectively competing in the market. Counting the number of firms in a market requires delineating market boundaries. This exercise can be contentious, for instance in anti-trust litigation (Carlton, 2007). For this reason, it is important to estimate the number of firms rather than taking it as given within geographic or temporal boundaries.

2.1 Preliminaries

Producers

The economy is composed by V villages. In each village, there are measure one homogeneous producers, each producing a unit of output and there are n buyers who compete for these producers’ output, the monopsony case being $n = 1$. The inverse supply buyer i faces in a village is:⁶

$$p_i = \alpha + \beta q_i + \gamma \sum_{j \neq i} q_j. \quad (1)$$

⁵To be clear, the buyer differentiation we discuss is distinct from differentiation of products across sellers.

⁶This inverse supply can be microfounded by assuming a representative producer whose cost function features love for variety. Specifically, the producer profit function is: $V(p_1, \dots, p_n, q_1, \dots, q_n) = q_0 + \sum_{i=1}^n p_i q_i - C(q_1, \dots, q_n) = q_0 + \sum_{i=1}^n p_i q_i - (\alpha \sum_{i=1}^n q_i + \frac{1}{2} \beta \sum_{i=1}^n q_i^2 + \gamma \sum_{j \neq i} q_i q_j)$, where q_0 is the output that is not sold to buyers (e.g., consumed, not harvested), p_i is the price paid by buyer i and q_i is the output sold to buyer i (the solution presented in this section assumes $q_0 > 0$). A representative agent strategy featuring love for variety may itself be considered a “reduced form” approach that aggregates heterogeneous producers having idiosyncratic preferences for each buyer.

This equation adapts the standard model of linear demand and differentiated producers (see, e.g., Vives, 2001) to a setting that features imperfect competition among buyers rather than sellers.

From the producers' perspective, buyers are differentiated at rate $\Gamma \equiv 1 - \frac{\gamma}{\beta}$. If $\Gamma = 0$, buyers are homogeneous: the slope of the inverse supply to own quantity equals the slope to a competitor's quantity. If $\Gamma = 1$, buyers are local monopsonists: a buyer's inverse supply does not depend on other buyers' quantities.

Buyers

The profit of buyer i in a village is given by

$$\pi_i = q_i(v_i - p_i), \tag{2}$$

where q_i is the quantity purchased, v_i is the wholesale price net of costs, and p_i is the (effective) price the buyer pays to producers. A given buyer can compete in multiple villages. However, the choices made across villages are assumed independent, allowing omission of a village index from the equations above.

Assume buyers are ex-ante symmetric in the resale price v . The experiment introduces a subsidy, s , for a share μ of the buyers, who then have a higher resale price. Therefore, $v_i = v + s$ for treatment buyers and $v_i = v$ for control buyers. Below we refer to variables for treatment (control) buyers with subscript T (C).

2.2 Industry Equilibrium

Assume conduct takes the form of Nash-in-quantities competition, in which each buyer sets q_i strategically, taking into account competitors' choices q_{-i} and market structure n . This equilibrium concept includes Cournot oligopsony and monopsony for $n = 1$.

Consider a group-symmetric equilibrium in which firms in the same treatment group behave similarly. The first order conditions associated with each group's profit maximization are

$$\begin{aligned} q_T &= \frac{(2\beta - \gamma)(v - \alpha) + s(2\beta + \gamma(n - 1)) - sv\gamma\mu n}{(2\beta - \gamma)(2\beta + \gamma(n - 1))}; \\ q_C &= \frac{(2\beta - \gamma)(v - \alpha) - s\gamma\mu n}{(2\beta - \gamma)(2\beta + \gamma(n - 1))}. \end{aligned} \tag{3}$$

By the inverse supply functions in Equation 1 equilibrium prices are

$$\begin{aligned}
p_T &= \frac{\alpha\beta(2\beta - \gamma) + v(2\beta - \gamma)(\beta + \gamma(n - 1)) + \beta\gamma\mu ns + s(\beta - \gamma)(2\beta + \gamma(n - 1))}{(2\beta - \gamma)(2\beta + \gamma(n - 1))}; \\
p_C &= \frac{(2\beta - \gamma)(\beta(\alpha + v) + \gamma v(n - 1)) + \beta\gamma\mu ns}{(2\beta - \gamma)(2\beta + \gamma(n - 1))}.
\end{aligned} \tag{4}$$

These prices imply variable markdowns, $\frac{p_i}{v_i}$. Treatment (control) quantities are increasing (decreasing) in the subsidy amount s and both are decreasing in the share of treated buyers μ . Both control and treatment prices are increasing in both s and μ . These intuitive comparative statics show how treatment changes the behavior of control firms, as well as treatment ones.

2.3 Empirical Identification of the Market Structure Parameters

We now review the equilibrium equations that identify the two parameters of the supply curve, Γ and β , and the effective number of firms, n .

2.3.1 Average Treatment-Control Differences in Quantities and Prices

For a given subsidy level, s , the differentiation parameter Γ and the slope of the supply curve β are identified by the treatment and control differences in prices and quantities

$$\Delta p \equiv p_T - p_C = \frac{s(\gamma - \beta)}{\gamma - 2\beta} = \frac{s\Gamma}{1 + \Gamma}; \tag{5}$$

$$\Delta q \equiv q_T - q_C = \frac{s}{2\beta - \gamma} = \frac{s}{\beta(1 + \Gamma)}. \tag{6}$$

Intuitively, Δp is increasing in Γ . If traders are homogeneous (i.e. $\Gamma = 0$), there can be only one price in the market; with higher differentiation, different prices can coexist and control traders can pay producers a price lower than do treated traders. Conversely, Δq is instead decreasing in Γ : if traders are more differentiated, market stealing from control to treatment will be smaller. Intuitively, Δq is also decreasing in β the slope of the inverse supply curve. These results do not depend on n or the share of treatment traders μ .

From the estimates of $\Gamma = 1 - \frac{\gamma}{\beta}$ and β , we can obtain the own-price price elasticity, η_{ii} , and the cross-price elasticities, η_{ij} :

$$\eta_{ii} \equiv \frac{\partial q_i}{\partial p_i} \frac{p_i}{q_i} = \frac{\beta + \gamma(n - 2)}{(\beta + \gamma(n - 1))(\beta - \gamma)} \frac{p_i}{q_i}; \tag{7}$$

$$\eta_{ji} \equiv \frac{\partial q_j}{\partial p_i} \frac{p_i}{q_j} = -\frac{\gamma}{(\beta + \gamma(n - 1))(\beta - \gamma)} \frac{p_i}{q_j}. \tag{8}$$

These results demonstrate how average individual-level treatment and control differences from a randomized experiment can be used to measure the extent to which individuals in a market are differentiated, and how this information can be used to back out the supply elasticity. These elasticities shape the quantity responses to price changes (including to changes induced by the subsidy) for a given number of competitors.

One approach to recover these elasticities would be to define the number of buyers n as the ones that are found to operate in a village, but as discussed, this would require delineating market boundaries. It is possible however to use additional information from the data to estimate market structure, rather than to take it as given. We discuss this approach in the next section.

2.3.2 Pass-through of the Wholesale Price

The model establishes a method to estimate the effective number of firms n , without the need to delineate market boundaries. For a given estimate of Γ , n is identified by the pass-through rate, which describes how producer prices p_i respond to a change in v , the component wholesale price that is common to all buyers. The pass-through rate is given by

$$\rho \equiv \frac{\partial p_C}{\partial v} = \frac{\partial p_T}{\partial v} = 1 - \frac{1}{1 + \Gamma + n(1 - \Gamma)}, \quad (9)$$

which is decreasing in Γ and increasing in n . In practice, Γ and n are estimated jointly, not sequentially.

3 Experimental Subsidies to Farm-Gate Buyers in the Sierra Leone Cocoa Industry

Setting. Cocoa is an important crop for Sierra Leone, where it is the largest agricultural export by value, comprising 8.6% of exports in 2017 according to the United Nations (2020). Approximately 75 percent of global supply of cocoa originates from West Africa (ICCO, 2019).

A brief schematic of the industry follows. Wholesalers are based in three main towns along the Moa river plain in Eastern Province, as shown in Figure 1. They source cocoa from a network of traders (intermediaries) with whom they typically have exclusive relations such that a trader almost always delivers cocoa only to a single wholesaler. Traders purchase cocoa from farmers in villages near the towns and deliver to wholesalers, who sell onwards to exporters in the provincial capital of Kenema or the national capital Freetown. Though Sierra Leone supplies only 0.77% of global cocoa exports, the structure of the country’s industry is similar to other exporters in West Africa,

though in Sierra Leone there is no government involvement in downstream purchases.

The provision of loans by traders to farmers is an important feature of this industry, suggesting the hypothesis that traders are potentially differentiated, lending them market power. Loans are typically given in the form of advance payment, when traders pay for cocoa in advance of delivery. Farmers use the advance payments for production or for consumption smoothing. Production mainly involves hiring workers to harvest the cocoa from the trees, which are planted infrequently. Farmers then pay interest on these advances by selling at a below market price for subsequent sales. Verbal contracts define the amount to be deducted from the final payment. Contracts may be enforced by customary legal institutions (see, e.g., Acemoglu et al., 2014 and Sandefur and Siddiqi, 2013), or through relational contracts, in which the farmers' fear of disrupting future business with the trader could be sufficient to avoid default (see, e.g., Fafchamps, 2003, Macchiavello and Morjaria, 2015, Blouin and Macchiavello, 2019).

SUTVA and experimental analysis within a competitive market. In this setting, competitive forces may cause a violation of the stable unit treatment value assumption (SUTVA) required for standard experimental analysis. Average differences in outcomes between treatment and control groups therefore cannot be interpreted as treatment effects in the potential outcomes framework of Rubin (1974).

In our setting, the SUTVA violation is plausible because treatment and control traders operate in the same village. If a village's supply of cocoa is not perfectly elastic, the offer of a higher price by treatment traders would induce a strategic response from control traders. The SUTVA could also be violated if randomization is conducted at the village level, for instance if the experiment offered the subsidy to all traders in certain treatment villages, while offering it to none in control villages. Each trader operates in 4.6 villages on average, with significant multimarket contact between traders. Traders' mobility (on motorcycles) and the relatively small geographic region of the cocoa producing area imply that, for a given trader, most markets are contestable. This becomes clear inspecting Figure 1 which shows the locations of the three towns Segbwema, Pendembu, and Kailahun where the five wholesalers operate, as well as the study villages from which study traders procure cocoa. All three towns and villages lie within 40 miles of each other.

While experimental studies in development economics often assume the SUTVA holds across villages, this assumption may be controversial, if not implausible. The model in Section 2 establishes a method to overcome this challenge. It provides a means to interpret the average differences between treatment and control prices and quantities when SUTVA does not hold. Further, it

demonstrates how combining these differences with an estimate of the pass-through rate identifies the number of firms competing in the market.

Experiment Design and Implementation. As with many export products, a key policy concern in the cocoa industry is how to upgrade average quality. The transmission of a quality price premium to farmers is a necessary condition to do this in a decentralized manner. Our experiment sought to demonstrate how changes in traders' marginal revenue for high quality cocoa translates into price signals received by cocoa farmers. We developed the experiment in partnership with five private wholesalers. Ultimately, the experimental sample comprised 80 traders, henceforth study traders.

When studying prices, it is important to focus on narrowly defined homogeneous goods, lest price differences reflect differences in quality. The quality of cocoa is heterogeneous, and market prices depend on a variety of characteristics including moisture content, mold, germination, lack of fermentation and a discoloration known as slate. Though there is no official standard of quality in the market, wholesalers and traders agree on broad determinants of quality that are consistent with international standards (see Fold, 2005). To implement the experiment, we developed with wholesalers a quality index that correlates well with baseline prices, described in further detail in Appendix A. When traders arrive at the warehouse, inspectors hired by the research team sampled 50 beans from each bag, and scored each bean to create an index of quality—grades A, B or C—that was applied to each bag. Formal grading was explained to traders as an initiative to make wholesaler pricing, which was already based to some extent on inspection of quality, more rigorous. The analysis in this paper focuses on grade A cocoa, the grade targeted by the experimental subsidy, unless otherwise specified.

The experiment was implemented as follows. From mid-October to the end of December of 2011, roughly the end of the harvest season, a randomly-selected subset of 40 traders we offered a bonus of 150 Leones per pound of cocoa sold —5.6% of the average wholesale price —when selling good quality (grade A) cocoa to the wholesalers. At the beginning of the experiment, traders were informed the bonus was because of increased demand for high quality cocoa. Randomization of the bonus treatment occurred at the individual trader level. We implement a pairwise randomization strategy (Bruhn and McKenzie, 2009), which matches traders within wholesalers according to their self-reported estimate of the volume of purchases since the beginning of the cocoa season and then assigned treatment and control within pairs. Of the 84 traders identified by wholesalers, four were outliers with respect to baseline quantity relative to other traders within the same wholesaler, and

could not be matched to other traders in our randomization strategy. Thus, the final sample selected for randomization was 80 traders.

3.1 Data and Summary Statistics

Over the course of the experiment, we collected a variety of original data from cocoa traders using three instruments: (i) a trader baseline survey, which recorded basic information on the trader and his business (all are male); (ii) a transaction survey, which for each transaction recorded the unit value paid to the farmer, the shipment weight in kilograms, and cocoa quality according to the grading scheme; and (iii) a farmer listing, administered at baseline and then in two follow-up rounds, in which traders were asked to list the farmers they buy from, and whether they had provided them with advance payment (in the last 12 months at baseline and in the last month in the two follow-ups). Appendix Figure F.1 shows a time line of the harvest season, indicating the times at which each instrument was deployed. Transaction data collection ran from September 24th, 2011 to December 31st, 2011 and the cocoa transaction survey was administered continuously over this period. We began paying bonuses to treatment traders on October 15th, 2011 until December 31st, 2011. The exact dates each trader responded to the transaction survey varies, given that they arrived at the wholesalers' warehouses at different times. Data collection was suspended for approximately two weeks and half between late November and early December because of project budget constraints due to a higher volume of recorded transactions than we had initially budgeted. We cannot be certain how this unexpected break in ability to pay affected trust in the subsidy, but we do observe that treatment traders continued to bring more quantity than control traders after the subsidy was reintroduced. The farmer listings were given to traders the first time they arrived in October, December, and January.

Our key outcome variables are the transaction price, measured as the unit value of each shipment, and a dummy for whether the trader had provided an advance payment to a farmer in the previous month, in either the second or third round of the farmer listing. The cocoa transaction survey was administered as follows. During the experiment, when traders arrived at the warehouse, inspectors from the research team measured quantity and quality of their shipment. Enumerators then asked traders the price per pound they paid to farmers and the name of the village where the cocoa mostly originated. Traders often mix cocoa from different farmers in the same bag, and so farmer prices reported are the average per unit purchase price paid by a trader for the cocoa in the bag.

Panel A of Table 1 presents summary statistics from the trader baseline survey and the first trader listing. Treatment and control groups are balanced on these trader-level covariates. In the

baseline listing, traders report purchasing cocoa from 123 villages. The average trader operates in 4.6 villages, and buys from 6 farmers per village. On average, based on the trader survey, there are 7.8 traders operating in a village. However, only 3.2 of these are study traders, suggesting that about 60% of the traders in the market are non-study traders (i.e. working with other wholesalers). Study traders report having given at least one loan to about 70% of the suppliers listed at baseline in the previous year.

Attrition. In the three weeks preceding the intervention, 60 of the 80 traders included in the study visited the warehouses (29 control and 31 treatment). Panel B of Table 1 reports balance for prices and quantities in the three weeks before the intervention, and shows there is a marginally significant difference in quantities ($p=.136$). During the experiment, 75 traders visited the warehouse (37 controls and 38 treatment). We include in all regressions randomization pair fixed effects and thus we effectively drop pairs including those traders that did not visit the warehouses.

4 Experimental and Quasi-Experimental Results

In this section we report the average differences between treatment and control traders in prices paid to farmers, provision of advance payments, and quantities of cocoa purchased, and the quasi-experimental results on pass-through.

4.1 Average Treatment and Control Differences

4.1.1 Farmer Prices

First, we examine treatment and control differences in prices (i.e., unit values) that traders pay to farmers. Enumerators asked traders the unit value for each shipment, which dividing by weight yields a measure of price. If the trader made payments at different times (e.g., before and after harvesting), enumerators recorded the total price paid to the farmer, including anything paid before harvest. We denote this price with \tilde{p} , so to differentiate it from the effective price, p , which accounts for the farmer’s valuation of advance payment and is defined formally in Section 5.1.1.

Figure 2 displays the price results graphically. It shows weekly averages for i) world prices,⁷ ii) wholesaler prices, iii) prices treatment traders paid to farmers, and iv) prices control traders paid to farmers. The vertical red line marks the inception of the experimental treatment. There are two

⁷Specifically, we report the front month price on the Intercontinental Exchange for the physical delivery of 10MT of exchange-grade cocoa from a variety of African, Asian and Central and South American origins to any of five U.S. delivery ports, with trading code NYCC. We convert prices from USD/metric ton to Leones/pound using the prevailing exchange rate of 1 USD=4,400 Leones.

observations to make. First, farmer prices follow closely wholesaler prices, which move with world prices. In particular, domestic prices respond to the sharp decrease in the world price that occurred in November 2011. Second, there is no obvious gap between the average prices that treatment and control traders pay to the farmers, either before or after the treatment begins. Suggestive evidence in line with the hypothesis that the SUTVA is violated comes from the fact that the treatment-control price gap is larger in the first weeks of the experiment (i.e., in the first three weeks of the experiment, the treatment coefficient is 31, $p=0.09$), and then decreases as control traders respond to higher prices offered by treatment traders.

We estimate the average treatment and control difference in prices using the following regression, where an observation is a shipment s delivered by trader i of randomization pair z in week t :

$$\tilde{p}_{sit} = \eta_{z(i)} + \eta_t + \beta_{\tilde{p}} \text{Treat}_i + \epsilon_{sit}, \quad (10)$$

where $\eta_{z(i)}$ and η_t are randomization pair and week fixed effects, respectively. We cluster standard errors at the unit of treatment (i.e., the trader). Results are similar when allowing for double clustering by trader and village (Cameron, Gelbach and Miller, 2012).

The term $\beta_{\tilde{p}}$ is the coefficient of interest. In Table 2, Column 1 presents the regression without week fixed effects, and $\hat{\beta}_{\tilde{p}} = -32.5$ (s.e.=47.2). Column 2 includes week fixed effects and yields $\hat{\beta}_{\tilde{p}} = -5.5$ (s.e. = 14.9). While the two coefficient estimates are not statistically distinguishable from each other or from zero, the coefficient is greater in absolute value without week effects, suggesting that selection in when to sell matters. It appears that the experiment induced treatment traders to stay longer in the market at the end of the season, when prices were lower.

The treatment may also have induced selection into which traders make purchases and the locations traders visit. To account for that possibility, columns 3 to 5 include controls referring alternatively to the trader, the village where the majority of cocoa in the shipment originated⁸, and both. See the notes to Table 2 for complete list of these controls. The coefficient $\hat{\beta}_{\tilde{p}}$ is quite stable across these columns, suggesting that the selection described above does not drive the results. Overall, across alternative specifications average prices paid to farmers are not different between treatment and control traders.

In Appendix Table F.1, we also test for effects on prices of B and C grade cocoa, which were not subsidized. We find a statistically significant difference for grade B prices, however the value is still far from the value of the subsidy. Field interviews suggest that treatment traders were somewhat

⁸Eighty of the 123 villages listed at baseline appear as “main village” in at least one shipment, covering approximately 85% of the suppliers listed at baseline.

more willing to pay the grade A price for cocoa that had some probability of being grade A.

4.1.2 Advance Payments

Second, we estimate average treatment-control differences in the provision of advance payments during the intervention period, using the following linear probability model,

$$AdvancePayment_{fi} = \eta_{z(i)} + \beta_a Treat_i + \nu_{fi}. \quad (11)$$

An observation is a farmer listed as a regular supplier in the trader baseline survey. $AdvancePayment_{fi}$ is an indicator of whether trader i reported paying farmer f an advance payment during the course of the experiment, in either of the two follow-up listing exercises.⁹

Table 3 presents estimates of the coefficient of interest, β_a . Column 1 presents the results of estimating Equation 11, which yields $\beta_a = 0.12$ (s.e. = 0.03), implying farmers reported by treatment traders in the baseline listing are 12 percentage points more likely to receive credit from these traders during the experimental period, relative to a control mean of 15 percent. In Columns 2 to 4, the coefficient does not change when adding trader controls, village controls, and both set of controls together.

Note that this advance provision may cover sales of cocoa for various grades, not just grade A. In Section 5.1.1, where we value advanced payments, we consider the implications of this issue.

4.1.3 Quantities

Third, we estimate average treatment-control differences in quantity purchased from farmers. Figure 3 shows the weekly amount purchased by the study traders together and then by treatment and control groups separately. Several patterns emerge. First, purchases of treatment and control are balanced in the two weeks before the intervention, while control quantities are higher three weeks before the beginning of the intervention. Second, throughout the intervention, treatment traders purchase substantially higher volumes than control ones. Third, total quantity purchased by study traders continues to increase after the beginning of the experiment. This observation is consistent with the idea that treatment traders gained market shares at the expense of non-study traders, as

⁹In the listing, we recorded data only on regular suppliers and it is not clear in which direction this selection may bias our estimates of the advance payment treatment-control difference: traders may be less likely to extend advances to irregular suppliers or, on the contrary, they may be using advances particularly to attract irregular suppliers. However, in the price regressions just reviewed, which include purchases for all suppliers, not just regular suppliers, average differences between treatment and control did not vary when controlling for the number of regular suppliers in the village. This provides some assurance that traders did not contract with regular suppliers differently from how they passed value to other farmers.

well as control traders. Fourth, toward the end of the experiment, there is a stark reduction in total quantities purchased, consistent with the season ending at that time.

These results are quantified more precisely using the regression model

$$q_{it} = \eta_{z(i)} + \eta_t + \beta_q \text{Treat}_i + \zeta_{it} \quad (12)$$

where an observation captures the total purchases of cocoa trader i in week t (including zeros). Table 4 presents estimates of the coefficient of interest, β_q . Column 1 presents the results of estimating Equation 12, which yields $\beta_q = 398.4$ (s.e. = 38.0), indicating that during the experiment treatment traders on average purchase 398 pounds per week more than control traders, or 349% more than the control mean.¹⁰ The results are robust when including trader controls in Column (2). Overall, this is a large impact of the treatment.

4.2 Quasi-Experimental Results: Pass-Through of the Wholesale Price

In this section, we estimate the pass-through rate to farmers of a common change in the wholesaler price (i.e., ρ , as defined in Equation 9). Estimation of ρ is complicated by the possibility of reverse causality, wherein local shocks to farmer costs and supply may affect farmer prices, which in turn affect wholesaler prices. To address this concern, we instrument wholesaler prices with the international price of cocoa, as reported by the Intercontinental Exchange. Given that Sierra Leone has a small share of the global production, it is plausible that changes in international prices are exogenous to supply conditions in Sierra Leone. Local prices are also highly correlated with international prices in the time series. Recall Figure 2, discussed in Section 4.1.1, which showed a stark reduction in prices paid to farmers (around 22%) in the final month of the experiment, following a reduction in wholesaler prices, and a decline in the world price.

Table 5 presents the results on pass-through from wholesale prices to farmer prices in a regression framework. In Column 1, we report estimates from two-stage least squares estimation, controlling for trader fixed effects and clustering by trader and date. The instrument has a very strong first stage, with the Kleibergen-Paap F Statistic being equal to 57.3. The coefficient estimate is $\rho = 0.91$ (s.e. = 0.06). Column 2 shows that the coefficient estimate is robust to including village fixed effects. In Column 3, we collapse data by date and run a time-series regression with standard errors

¹⁰Consistent with the large difference in quantities purchased, treatment traders were more than three times as likely to visit the warehouse during the experimental period than control ones. Throughout the experiment, we did not receive any complaint from either wholesalers or traders suggesting that control traders were switching to different wholesalers. This is consistent with the fact that the experiment did not change the wholesaler price for control traders.

robust heteroskedasticity and autocorrelation up to 70 lags (HAC Newey-West). The coefficient falls slightly, but is still high at $\rho = 0.85$ (s.e. = 0.03). Overall, these results suggest almost all of changes in the wholesale price are passed through to farmers. These results are consistent with the findings of Gayi and Tsowou (2015), who show that cocoa farmer prices in several West African countries have been very responsive to world prices in the last two decades, with a pass-through rate of around 90%.

One caveat to these results is that this estimate of the pass-through rate does not include pass-through of value to farmers on non-price margins. Given the relatively low frequency of international price movements, measuring pass-through in these terms would require collecting data on advance payments over multiple seasons. Therefore, the estimated pass-through rate, which is already quite high, is likely to be a lower bound of the pass-through one would measure when accounting for the value of advanced payment.¹¹

5 Estimating the Model Parameters

In this section, we use the experimental and quasi-experimental results to estimate the model parameters. The analysis proceeds in five steps. First, we recover trader differentiation, Γ , and supply slope, β from our experimental results. Second, we combine Γ with our quasi-experimental estimate of the pass-through rate to recover n . Third, we combine β , Γ , and n , to calculate own-price and cross-price supply elasticities, and quantify the price and quantity response to the experiment relative to a counterfactual without the experiment. Fourth, we discuss key assumptions of our approach and, where possible, validate these assumptions using the data. Fifth, we run simple counterfactual exercises to illustrate the importance of estimating all the market structure parameters.

5.1 Estimating Trader Differentiation (Γ) and Supply Slope (β) from the Experiment

To recover trader differentiation, Γ , and supply slope, β , we match the moments in Equations (5) and (6) to their empirical analogs from the experiment. Here, we first discuss how to combine data on prices and advance payments into an effective price that summarizes the present value of

¹¹Advance payment is the only margin of non-price competition on which we have data. Another potential margin could be that traders provide price insurance to farmers. Additional analysis shows the pass-through rate does not vary between treatment and control (p-value=0.43). This suggests that treatment traders do not provide (additional) price insurance to farmers relative to control ones. This finding supports our approach in Section 5.1.1, where we calculate the effective price using only price and advance payment provision.

the transaction for the farmer. We then present the main results from the estimation, substituting this effective prices for the price p_i in Equation 1 of our model.

5.1.1 Preliminary Step: Treatment-Control Differences in Effective Prices

The effective price paid to the farmer, p_i can be written as

$$p_i = \frac{\tilde{p}_i}{1 - \sigma_i(1 - \lambda)} \left(1 - \sigma_i \left(1 - \frac{\lambda}{\delta}\right)\right), \quad (13)$$

where \tilde{p}_i is the price paid by trader i , inferred from unit transaction values as described in Section 4.1.1; σ_i is the share transactions of trader i using advance payments; λ is the price discount accounting for interest (i.e., $\lambda = \frac{1}{1+r} < 1$); and δ is the farmer’s discount factor, which captures the subjective value advance payment. Appendix B.1 derives the equation.¹²

To obtain p_i , we require values for λ and δ . We recover these from the data, starting with a simple approach that uses cross-sectional variation in price and advance payment provision. We then show that quantitative results of our estimation are similar using different approaches that leverage different sources of variation, including an auxiliary experiment that measures farmers’ subjective discount factor.

Baseline Approach: Cross-Sectional Variation in Farmer Prices and Advance Payment Provision. As a starting point, we assume that the discount factor is the same for farmers and traders (i.e., $\delta = \lambda \Rightarrow p_i = \frac{\tilde{p}_i}{1 - \sigma_i(1 - \lambda)}$). In the first approach, we infer $\delta = \lambda$ from the baseline cross sectional relationship between prices and advance payments. Since we observe payment amounts at the village level but not at the transaction level, our focus is on village-level average prices and on the share of farmers receiving advance payments in the village.¹³ Appendix Table B.1 reports the results of a regression of the price on the share of advance payments. Moving from a village where no farmer receives advance payments at baseline to a village where each farmer re-

¹²We abstract from dynamic features of the farmer-trader relationship that may support the repayment of the advance. Accounting for these elements would require a repeated game framework, featuring multiple choice variables for the traders, each dependent on market structure, which does not lend itself easily to closed-form solutions for the treatment-control differences, substantially complicating estimation. We assume instead that traders face a separable problem. First, they set their effective price conditional on the inverse supply curve and competition they face. Second, for a given effective price, they choose the combination of payments to be made at different times. We do not model this second step. When making their sale choices, farmers consider the effective price, not its composition. As a result the model captures a continuum of potential equilibrium contracts.

¹³Village-level averages come from aggregating traders’ baseline responses on prices, locations of activity, and number of suppliers. Here, we use villages as spatial unit to study the relationship between prices and advance payments. This is not inconsistent with our later discussion that villages may not be the relevant definition of market size. Our goal here is to estimate the slope of total payments with respect to advance payment provision. This requires partitioning farmers and using the partitions as data points. Villages are one of the many possible partitions, but a natural one to use (among other reasons, because we have covariates at the village level).

ceiving advance payments decreases shipment prices paid by the trader by approximately 150 Leones from an average of 3,138 , so $\delta = \lambda = 0.95$.

Having calibrated λ and δ , it is possible to compute the average treatment-control difference in effective prices, $p_T - p_C$. Table 6—Panel A summarizes the results (Columns 1-3). Average prices for control and treatment traders are: $\tilde{p}_C = 2,987$, $\tilde{p}_T = 2,982$ (from Table 2, Column 2), $\sigma_C = 0.15$, $\sigma_T = 0.27$ (from Table 3, Column 1). With these values, the average effective prices implied by Equation (13) are $p_C = 3,010$ and $p_T = 3,022.8$ and the average treatment and control difference in effective prices is 12.8 Leones, with 90% bootstrapped confidence intervals [-17.8,47.9].¹⁴

Sensitivity to Alternative Approaches to Valuing Advance Payments. The baseline result is robust to alternative approaches to estimating λ and δ . Appendix B.2 presents details of the alternative strategies. Here, we provide a brief discussion. Table 6—Panel B presents the results in Columns 1-3. First, we assume again $\delta = \lambda$ and infer the value of advance payments from the covariance between treatment and control differences in prices and in advance payments. The slope between the two response margins identifies their relative value, or how much less a trader who increases his advance payments needs to adjust his prices. The estimation suggests that when treatment traders provide credit but control traders do not, shipment prices would fall by 221 Leones, from a baseline of 3,138 Leones. Thus, $\lambda = 0.93$ and the difference in effective prices between treatment and control is 20.3 Leones.

Second, we use the interest rate offered by Rural and Agricultural Banks to calibrate λ . As a lower bound for this value, we consider the rate of 2% per month. Assuming a loan duration of a month, i.e. approximately half of the duration of the intervention, $\lambda = 0.98$ (higher values of the interest rate would give values of λ closer to our baseline estimate). In this case the difference in effective prices between treatment and control traders is 1.7 Leones.

Third, we allow the interest rate to differ from the subjective rate at which farmers’ value future advances, or $\delta \neq \lambda$. To measure the farmer’s discount factor, in November 2020 we conducted an additional incentivized lab-in-the-field experiment in three of the villages included in the main field experiment and during the same season. We asked farmers to make a number of binary choices between receiving money today or in the future. We estimate a median monthly subjective discount factor of $\delta = 0.914$. This rate is in line with recent experimental evidence on time preferences from other African countries (e.g., Balakrishnan, Haushofer and Jakiela, 2020). When we combine this estimate of δ with our baseline estimate of λ from cross-sectional variation in prices and credit (i.e.,

¹⁴We treat the auxiliary parameter λ as a calibrated parameters and we do not account for its sampling variance when computing bootstrapped standard errors. Instead, we present sensitivity of our results to alternative methods to calibrate λ .

$\lambda = 0.95$), the difference in effective prices between treatment and control traders is 27.2 Leones.

Fourth, we modify our baseline approach to deal with the observation that some of the advance payments may have been given with the expectation that a farmer would deliver non-grade A cocoa (as per our discussion in Section 4.1.2). For this purpose, we scale down both the baseline village-level credit share and the credit shares in each treatment group during the experiment by the share of grade A transactions. We obtain $\lambda = 0.92$, $\sigma_C = 0.068$, $\sigma_T = 0.17$. In turn, the implied difference in effective prices between treatment and control traders is 19.6 Leones [-11,54.7]. Through a similar procedure, the difference in effective prices for grades B and C is shown to be 31.1 when pooled together, and 37.8 when restricting to grade B only. These values are within the 90% confidence interval of the treatment-control difference in effective prices for grade A.

5.1.2 Estimating Γ and β

We can now estimate Γ and β by matching the moments in Equations (5) and (6) to their empirical analogs, using the difference in *effective* prices in Equation (5). Intuitively, we recover Γ (the degree of differentiation among traders) from the price treatment-control difference and β (the slope of the trader’s inverse weekly supply) from the quantity difference.

Columns 4 and 5 of Table 6 report the result. The first row shows results with our baseline approach to compute effective prices ($\lambda = \delta = 0.95$). We obtain $\Gamma = 0.09$ (90% C.I. [-.11,.47]) and $\beta = .35$ (90% C.I. [.24,.48]).¹⁵

The remaining rows in Table 6 illustrate the sensitivity of these estimates to alternative approaches to valuing advance payments. Estimates of trader differentiation Γ range between .01 and .22, with the highest value when calibrating the farmer’s discount factor from the auxiliary experiment, which also gives the highest difference in effective prices. This result illustrates how the value of credit is enhanced when credit markets are incomplete and $\delta < \lambda$. Estimates of the inverse supply slope β span between .32 and .37.

Overall, these results suggest that regardless of how advance payments are valued, traders appear fairly undifferentiated. This results is consistent with other empirical findings from sub-Saharan Africa (Fisman and Raturi, 2004; Ghani and Reed, 2021) and from China (Fabbri and Klapper, 2016), which show that buyers provide relatively more trade credit to suppliers with whom they have relatively less bargaining power. One explanation for this result could be the existence of customary legal enforcement. In this setting, disputes can be brought before a customary law court run by the Paramount Chief and his or her deputies, which has the authority to levy penalties for

¹⁵From Γ and β , it is then straightforward to obtain $\gamma = (1 - \Gamma) * \beta = .318$

breach of contractual obligations. The threat of enforcement in such courts may obviate the need for relational contracts requiring market power to sustain credit contracts. Another is that competition reduces holdup concerns on the lender side and it may thus increase borrowers' incentives to establish creditworthiness (as in Fisman and Raturi, 2004).

5.2 Estimating the Number of Competitors (n) from Quasi-Experimental Variation in Cocoa World Prices

We use the quasi-experimental estimate of the pass-through rate Section 4.2 to estimate the number of buyers effectively competing for supply (Equation 9).

In the average village, we observed 7.8 traders operating, already perhaps a large number. Using our estimate of $\rho = 0.91$ from Table 5, and an estimate of Γ using the baseline approach to valuing advance payments delivers an estimate of $n = 11.1$ with 90% C.I [8.8,18.9] (first row of Table 6, Column 6). These estimates of n imply that, according to the model, traders behave as if the number of their competitors were about 40% percent higher than observed in the average village.

This result confirms the intuition that village markets operate as if they are highly contestable: the option to sell to other traders shapes competition, not the actual number of traders actually purchasing from each farmer. Sensitivity analysis in the remaining rows of Table 6 shows a range of estimates of n between 10.3 and 12.8, suggesting that the result is robust across alternative approaches to valuing advance payment.

We obtain similar results when using alternative moments from the model, in the spirit of an overidentification test. Specifically, we derive theoretical expressions for the *percent* difference in prices and quantities between treatment and control traders and match them to their empirical counterparts. This approach yields estimates which are quite similar; for instance, when using $\lambda = .95$, we obtain $\Gamma = 0.097$ and $n = 11.07$. Appendix Figures C.1 and C.2 also show that, in the case of model misspecification, estimation using absolute or percent differences would yield considerably different estimates. Appendix C provides details.

Consistent with the idea that the number of traders observed operating in a village do not necessarily measure the market structure, we do not detect statistically significant impacts of the number of treated traders in the village on the treatment-control differences in prices, advance payments, and quantities.

5.3 Supply Elasticities and the Response to the Experiment

With the model’s estimated parameters (Γ, β, n) in hand, one can calculate: i) own-price and cross-price elasticities (from Equations 7 and 8); and ii) the impact of the experiment, relative to a counterfactual without the experiment.¹⁶

5.3.1 Own- and cross-price supply elasticities

Using the baseline parameter estimates and average prices and weekly quantities, we obtain the own- and cross-price elasticities facing each trader. Note results are similar when using counterfactual prices and quantities in the absence of the experiment, which we derive next.

The own-price elasticity is very high, $\eta_{ii} = 324$. This is exactly what we would expect in a competitive market. A small increase in the prices of one of the competitors leads to a large increase in supply. In the limit case of perfect competition with atomistic firms, this value would be infinity. The cross-price elasticity is $\eta_{ji} = -32$, which implies that if one of the competitors increase the price by 1%, each of the other 10 competitors loses 32% of their supply. This is consistent with a very competitive market.

5.3.2 The Impact of the Experiment

Through the lens of the model, we can quantify the impact of the randomized subsidy on prices and quantities of control and treatment traders, relative to counterfactual prices and quantities in the absence of the experiment. Without SUTVA, the counterfactual is not observed directly in the experimental control group, and must be inferred from theory. The main results of this calculation are described here, while Appendix D provides the equations.

The analysis proceeds in two steps. First, we calculate the effect of the subsidy on effective prices paid by both treatment and control traders, relative to a counterfactual without the experiment. The derivatives of the prices of treatment and control traders with respect to the subsidy value are pinned down by Γ, n , and μ . Recall μ is the share of treatment traders in the market, or 0.2. Using results from Table 6 –Panel A, we find that, in response to a subsidy of 150 Leones per pound, control traders increased (effective) prices by 24.7 Leones and treatment traders by 37.5 (hence, a difference of 12.8 Leones). Appendix Figure F.2 shows, for the estimated values of the competition parameters, the increase in the treatment and control prices in response to a unit subsidy, relative to the scenario without the experiment, as a function of the share of treated traders, $\mu \in (0, 1)$.

¹⁶Note that if one were sure about market boundaries, one could count the number of firms within a relevant geographic market and use that as an input to compute elasticities.

Second, using the estimate of β , as well as Γ , n , and μ , we compute the quantity responses to the experiment. In response to the subsidy, treatment traders increased supply by 324 pounds (a 174% increase relative to a counterfactual without the experiment) and control traders reduced supply by 71 pounds (a 38% decrease). This makes sense when recalling that there are four non-treated traders for each treated trader. Aggregate supply increased only by 4.1% and 88% of the increase in quantity for treatment traders comes from market stealing, a result reminiscent of Rotemberg (2019). A priori, one might have expected the majority of the increase to have come from market stealing, given that the experiment was implemented at harvest time and farmers had limited options to increase their supply in response to the price changes (e.g. reducing processing losses). By this time, production capacity is fixed by the number of trees that have been planted.

5.4 Discussion of Assumptions and Model Validation

The model presented in Section 2 makes a number of assumptions. To be transparent, we discuss these assumptions here and, where possible, provide additional evidence in their support.

Conduct. Our model assumes a specific equilibrium concept, Nash-in-quantities. Here we provide evidence against alternative hypotheses. An alternative model is a segmented monopsony, in which there are many traders, but each prices as a monopsonist facing a distinct set of farmers. The small difference in (effective) prices between treatment and control could reflect a scenario where each trader has a very high degree of market power but, given some alternative (i.e., non-linear) functional form of farmer supply, does not raise the price in response to the subsidy. The large quantity responses to treatment and the high degree of pass-through to common shocks provide initial evidence against this interpretation. Analysis of the implied elasticities provides additional evidence against segmented monopsony, under a *general* supply curve. A monopsonistic trader with general inverse supply curve $p(q)$ will price following Lerner’s condition: $p = v \frac{\epsilon}{\epsilon+1}$, where ϵ is the own price elasticity. Assuming this pricing condition holds and recalling that $v_C = v$, $v_T = v + s$, one recovers ϵ from the treatment-control difference in prices, $p_T - p_C = s \frac{\epsilon}{\epsilon+1}$. This gives a very small elasticity: $\epsilon = 0.093$. However, under segmented monopsony, one can also estimate ϵ from the ratio in the percent treatment-control differences in quantity and prices: $\frac{(q_T - q_C)/q_C}{(p_T - p_C)/p_C}$. In this case, given the large quantity difference, the estimated elasticity would be very high, $\epsilon = 812$. The inconsistency between these estimates contradicts the assumption of segmented monopsony.

A second alternative model is one in which traders operate in the same market, but have formed a cartel that can price as a monopsonist. This model is inconsistent with our finding a high pass-through rate, and large implied number of firms in the market. Were traders pricing as a

cartel, the number of firms identified by the pass-through rate would be one. Setting aside the quasi-experimental evidence from the pass-through rate, the fact that treatment-control differences are small for prices could be consistent with treatment and control buyers forming a cartel to take advantage of the subsidy by passing quantity to treatment from control traders. Collusion is however inconsistent with the large differential response of treatment traders in terms of advance payment provision. In addition, we note that collusion of this form would require not just an agreement between a treatment and a control trader to game the incentive system, but also collaboration among treatment traders (since otherwise a non-colluding treatment trader could steal the suppliers of the treatment-control pair cartel). The latter is a more demanding form of collusion and it faces the standard enforcement problems of a cartel.

Replicating our main estimation procedure using Bertrand competition, while retaining other assumptions on producers and buyers, delivers unrealistic parameter values (i.e., a value of Γ larger than one). This suggests that quantity may be the relevant strategic choice variable in the setting. As it is well known, Cournot outcomes can also be interpreted as reduced-form outcomes for price competition with quantity constraints (Kreps and Scheinkman, 1983). Quantity constraints arising from only being able to carry so many bags on a truck or motorcycle are relevant in this setting.

One might also posit a model of monopsonistic competition (adapted from the more common monopolistic competition case): Dixit and Stiglitz (1977) predicts a markdown on the subsidy equal to the markdown observed in the baseline data; Ottaviano et al. (2002) predicts a difference between treatment and control traders of one-half of the subsidy value. Neither of these predictions finds support in the data. For instance, farmer prices are on average 92% of the wholesaler prices. Under constant markdown case, this would imply a difference in effective prices of at least 135 Leones between treatment and control traders.

Functional Forms. The model assumes linear supply, rather than working with an unrestricted supply elasticity. Among other reasons, the use of linear supply may be a concern because linear supply can be microfounded with a representative agent approach, but not with a discrete choice problem (see, e.g., Jaffe and Weyl, 2010; Armstrong and Vickers, 2015).

It is possible to examine sensitivity of the results to this assumption. Consider an alternative inverse supply with curvature ζ : $p_i = \alpha + \beta q_i^\zeta + \gamma \sum_{j \neq i} q_j^\zeta$ (in the baseline linear inverse supply of Equation 1, $\zeta = 1$). Using the first order conditions and the price and quantity levels and treatment-control differences as estimating moments, we can recover the vector of parameters $(\beta, \Gamma, \alpha, v)$ as a function of the curvature parameter ζ . This analysis suggests that our key result of low trader differentiation is robust to relaxing the assumption of linear supply. Our benchmark $\zeta = 1$

corresponds to $v = 3,050$, which is very close to the sum of the effective price in the control group (3,010) and the average transport costs (49, as we discuss in the next paragraph). At $\zeta = .9$, v is equal to 3,059. Higher values of v , which imply traders get a larger share of surplus, imply lower ζ and, crucially, lower Γ (for instance, the upper bound of v equal to the average wholesaler price, or 3,260 Leones per pound, corresponds to $\zeta = 0.29$ and $\Gamma = 0.03$). Higher values of ζ are unlikely because they would imply that traders systematically make losses. For instance, when $\zeta = 2$, $v=3,017$. Even in this extreme case, Γ is still low (0.19).

Another assumption is that estimated v and p do not depend on quantities, ruling out non-linear pricing and other non-constant trader marginal costs (see, e.g., Attanasio and Pastorino, 2020). We also assume that the only cost for the trader is the crop purchase cost and that this cost is linear. This is a reasonable approximation given that the crop purchase costs are around 90% of resale prices. However, traders do bear other costs, for instance to transport the crop. Our estimates of the differentiation rate are intuitively robust to the introduction of other constant marginal costs. Appendix Table F.2 shows that unit transport costs of treatment traders do change slightly in response to the subsidy (a reduction of approximately 14 Leones, from a control mean of 49 Leones). Columns (3) and (4) of the same table suggest that they are more likely to use a truck to transport the crop, instead of motorbikes. In a simple twist of the model, the reduction in unit costs for treatment traders has the same effect of an increase in the subsidy value (i.e., from 150 to 164 Leones). Accounting for the change in transport cost, our estimate of Γ would slightly decrease (from 0.092 to 0.085).¹⁷

Representative agents. Agents are assumed to be symmetric, aside from the heterogeneity introduced by the experiment (i.e., the experimental subsidy, s). Appendix E.1 shows how the model can be extended to include heterogeneity in traders' marginal revenue and a heterogeneous differentiation rate across different pairs of traders. In principle, one could estimate Γ separately in each location and then compute the average of the parameters across villages. In practice, for our specific experiment, estimating separate parameters in each location (using information on the main village of provenience of the cocoa in the shipment) delivers results that are too noisy to be useful.¹⁸ A related concern is the presence of non-study traders. These comprise about 60% of the traders operating in the study region and, in principle, they could be different from the study

¹⁷Explicitly modeling a non-linear (e.g. quadratic) cost would complicate the relationship between the treatment-control differences and the parameters of interest. One would need additional moments (e.g. higher-order powers of the treatment-control differences) to achieve identification.

¹⁸For instance, with few observations per village, the treatment-control difference in effective prices is often either negative or larger than 150 Leones (i.e., the subsidy value), which in both cases implies a negative value of Γ . However, it is reassuring that when we include village fixed effects in the regressions with pooled data, results are very similar to the ones presented in the text.

traders (control and treatment) at baseline. The model presented in Section 2 is robust to the presence of such traders.¹⁹

Timing. Our model is static while the interaction between farmers and traders unfolds over an entire season. The experiment only ran until the end of the harvest season. Traders and farmers may have behaved differently in a multi-season trial. Again, it is plausible to assume that in a longer experiment the degree of differentiation would be even lower. For instance, in a longer experiment, farmers may have been more willing to switch to other buyers. Future research could assess whether varying the duration of the experiment leads to substantially different results. With this caveat in mind, we however believe that running the experiment until the end of the season was a reasonable length. The subsequent harvest season follows seven-eight months of inactivity and new trading relationships may potentially arise during that period. In addition, the high pass-through rate we described in Section 4.2 suggests that traders respond to high-frequency price changes, which are likely to be more transitory than our experimental season-long subsidy. The fact that the experiment lasted until the end of the season also suggests that traders had enough time to learn about the subsidy of their competitors, in line with the assumption of perfect information in the model. We also observe that the variation in prices induced by our experiment is less transitory than daily or weekly price variation used in many studies of pass-through.

5.5 Counterfactual Experiments

Estimates of the differentiation rate and effective number of firms in the market allow for the analysis of the impacts of subsidies that target subsets of firms in the market, a feature of many industrial (agricultural) policies. When subsidies are offered only to a subset of firms, they have direct effects through changes in the prices paid by treated firms, and indirect effects through the strategic response of untreated firms. The model may be used to simulate the general equilibrium effect of the policy, for different shares of subsidized firms. For example, in our context, a government might wish to subsidize agricultural output by raising the wholesale price. The model facilitates different subsidy interventions in terms of their *return on investment*: the ratio of benefits in terms of incremental farmer revenues and costs, or total expenditure on the subsidy.

Figure 4 reports the return on investment in three counterfactual scenarios to illustrate the importance of the differentiation parameter in this policy analysis. In each of the figures, we vary Γ , while keeping the pass-through rate ρ constant at 0.91, and thus adjusting n according to

¹⁹Appendix E.2 presents an extension of the model where only a share σ of traders is included in the study, and thus study treatments are a share $\sigma\mu$ of traders. Non-study traders have a resale price, v' , that possibly differs from the study traders' one, v . The equilibrium treatment-control price difference Δp (Equation 5) is unchanged.

Equation 9. Panel A shows that the return on investment of an intervention that provides subsidies to one-fifth of the traders (similar to our experiment) is increasing on Γ , with returns more than eight times as high for Γ close to one than for Γ close to zero. Panel B shows that the additional benefits of subsidizing more traders ($\mu = .8$ vs $\mu = .2$) are decreasing in Γ : the high-intensity intervention gives 20% higher returns at low values of Γ , but similar returns at high values of Γ . In Panel C, we consider two types of traders, each comprising half of the traders, which differ in the resale value and thus in the equilibrium quantity purchased (at $\Gamma = .09$, the difference in quantity is about five-fold). We consider an intervention that subsidizes half of each type of traders and one that subsidizes all the large traders (and none of the small ones). The former has always higher returns, but its relative benefits are decreasing in Γ . These counterfactuals highlight the importance of estimating separately the two parameters shaping pass-through (i.e. Γ, n).

6 Conclusion

The potential outcomes framework for experimental analysis is not valid when treatment and control agents interact strategically, a feature of settings in which experimental subjects are participants in the same market.

We have shown that, when SUTVA fails, individual-level randomized subsidies can identify market structure parameters. The average difference in prices paid by treatment and control agents informs an intuitive test of the degree of differentiation between them: only if agents are differentiated can there be systematic differences in the average prices paid by subsidized and non-subsidized agents. Combining an estimate of differentiation with the pass-through rate reveals the number of firms competing in the market in a Nash-in-quantities equilibrium.

Overall, the evidence suggests the Sierra Leone farm-gate cocoa market is highly competitive. The pass-through rate is high, farm-gate traders exhibit a low degree of differentiation, and own-price and cross-price elasticities are high. However, while these findings are suggestive of an overall high degree of competition at the farm gate, firms at downstream levels the supply chain not studied here (i.e., wholesalers, exporters) may have substantially more pricing power. In Sierra Leone, where exportation is organized by the private sector, Figure 2 showed that, though wholesaler prices respond somewhat to changes in the international price, pass-through is lower at that stage of the value chain. Identifying whether this lower level of price pass-through may be explained by weak competition among wholesalers or exporters is an important area for future research.

References

- Acemoglu, Daron, Tristan Reed, and James A Robinson.** 2014. “Chiefs: Economic development and elite control of civil society in Sierra Leone.” *Journal of Political Economy*, 122(2): 319–368.
- Aker, Jenny C.** 2010. “Information from markets near and far: Mobile phones and agricultural markets in Niger.” *American Economic Journal: Applied Economics*, 2(3): 46–59.
- Andersen, Steffen, Glenn W Harrison, Morten I Lau, and E Elisabet Rutström.** 2008. “Eliciting risk and time preferences.” *Econometrica*, 76(3): 583–618.
- Antràs, Pol, and Arnaud Costinot.** 2011. “Intermediated trade.” *Quarterly Journal of Economics*, 126(3): 1319–1374.
- Antras, Pol, and C Fritz Foley.** 2015. “Poultry in motion: a study of international trade finance practices.” *Journal of Political Economy*, 123(4): 853–901.
- Armstrong, Mark, and John Vickers.** 2015. “Which demand systems can be generated by discrete choice?” *Journal of Economic Theory*, 158 293–307.
- Attanasio, Orazio, and Elena Pastorino.** 2020. “Nonlinear pricing in village economies.” *Econometrica*, 88(1): 207–263.
- Baird, Sarah, J Aislinn Bohren, Craig McIntosh, and Berk Özler.** 2018. “Optimal design of experiments in the presence of interference.” *Review of Economics and Statistics*, 100(5): 844–860.
- Balakrishnan, Uttara, Johannes Haushofer, and Pamela Jakiela.** 2020. “How soon is now? Evidence of present bias from convex time budget experiments.” *Experimental Economics*, 23(2): 294–321.
- Bardhan, Pranab K.** 1980. “Interlocking Factor Markets and Agrarian Development: A Review of Issues.” *Oxford Economic Papers*, 32(1): 82–98.
- Bardhan, Pranab, Dilip Mookherjee, and Masatoshi Tsumagari.** 2013. “Middlemen margins and globalization.” *American Economic Journal: Microeconomics*, 5(4): 81–119.
- Bell, Clive.** 1988. “Credit markets and interlinked transactions.” In *Handbook of Development Economics*. eds. by Chenery, Hollis, and T.N. Srinivasan, 1 of Handbook of Development Economics: Elsevier, , Chap. 16 763–830.
- Bergquist, Lauren Falcao, and Michael Dinerstein.** 2020. “Competition and entry in agricultural markets: Experimental evidence from Kenya.” *American Economic Review*, 110(12): 3705–47.
- Berry, Steven, Martin Gaynor, and Fiona Scott Morton.** 2019. “Do increasing markups matter? lessons from empirical industrial organization.” *Journal of Economic Perspectives*, 33(3): 44–68.

- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan.** 2004. “How Much Should We Trust Differences-in-Differences Estimates?” *Quarterly Journal of Economics*, 119(1): 249–275.
- Blouin, Arthur, and Rocco Macchiavello.** 2019. “Strategic default in the international coffee market.” *Quarterly Journal of Economics*, 134(2): 895–951.
- Breza, Emily, and Cynthia Kinnan.** 2016. “Measuring the Equilibrium Impacts of Credit: Evidence from the Indian Microfinance Crisis.” *Manuscript, Northwestern University*.
- Bruhn, Miriam, and David McKenzie.** 2009. “In pursuit of balance: Randomization in practice in development field experiments.” *American Economic Journal: Applied Economics*, 1(4): 200–232.
- Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel.** 2018. “Sell low and buy high: arbitrage and local price effects in Kenyan markets.” *Quarterly Journal of Economics*, 134(2): 785–842.
- Busso, Matias, and Sebastian Galiani.** 2019. “The causal effect of competition on prices and quality: Evidence from a field experiment.” *American Economic Journal: Applied Economics*, 11(1): 33–56.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller.** 2008. “Bootstrap-based improvements for inference with clustered errors.” *The Review of Economics and Statistics*, 90(3): 414–427.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller.** 2012. “Robust inference with multiway clustering.” *Journal of Business & Economic Statistics*.
- CAOBISCO/ECA/FCC.** 2015. “Cocoa Beans: Chocolate and Cocoa Industry Quality Requirements.” *Memorandum*.
- Carlton, Dennis W.** 2007. “Does antitrust need to be modernized?” *Journal of Economic Perspectives*, 21(3): 155–176.
- Casaburi, Lorenzo, Rachel Glennerster, Tavneet Suri, and Sullay Kamara.** 2014. “Providing collateral and improving product market access for smallholder farmers. A randomised evaluation of inventory credit in Sierra Leone.” *3ie Impact Evaluation Report*, 14.
- Casaburi, Lorenzo, and Rocco Macchiavello.** 2019. “Demand and supply of infrequent payments as a commitment device: evidence from Kenya.” *American Economic Review*, 109(2): 523–55.
- Casaburi, Lorenzo, and Jack Willis.** 2018. “Time versus state in insurance: Experimental evidence from contract farming in Kenya.” *American Economic Review*, 108(12): 3778–3813.
- Chatterjee, Shoumitro.** 2020. “Market Power and Spatial Competition in India.” Technical report, The Pennsylvania State University manuscript.
- Chatzouz, Moustafa, Áron Gereben, Frank Lang, and Wouter Torfs.** 2017. “Credit guarantee schemes for SME lending in Western Europe.” Technical report, EIF Working Paper.

- Chau, Nancy H, Hideaki Goto, and Ravi Kanbur.** 2016. “Middlemen, fair traders, and poverty.” *The Journal of Economic Inequality*, 14(1): 81–108.
- Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White.** 2020. “Measuring time preferences.” *Journal of Economic Literature*, 58(2): 299–347.
- Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora.** 2013. “Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment.” *Quarterly Journal of Economics*, 128(2): 531–580.
- Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran.** 2018. “The price effects of cash versus in-kind transfers.” *Review of Economic Studies*, 86(1): 240–281.
- David, Sonii.** 2005. “Learning about Sustainable Cocoa Production: A Guide for Participatory Farmer Training 1. Integrated Crop and Pest Management.” *Sustainable Tree Crops Program, International Institute of Tropical Agriculture, Yaounde, Cameroon.*
- Dillon, Brian, and Chelsey Dambro.** 2017. “How competitive are crop markets in Sub-Saharan Africa?” *American Journal of Agricultural Economics*, 99(5): 1344–1361.
- Dillon, Brian M, and Christopher B Barrett.** 2015. “Global oil prices and local food prices: Evidence from east africa.” *American Journal of Agricultural Economics*, p. aav040.
- Dixit, Avinash K, and Joseph E Stiglitz.** 1977. “Monopolistic competition and optimum product diversity.” *American Economic Review*, 67(3): 297–308.
- Donaldson, Dave.** 2015. “The Gains from Market Integration.” *Annual Review of Economics*, 7(1): 619–647.
- Druilhe, Zoe, and Jesus Barreiro-Hurle.** 2012. “Fertilizer subsidies in Sub-Saharan Africa..” Technical report, ESA Working Paper 12-04. Food and Agriculture Organization.
- Efron, Bradley.** 1981. “Nonparametric standard errors and confidence intervals.” *Canadian Journal of Statistics*, 9(2): 139–158.
- Emran, M Shahe, Dilip Mookherjee, Forhad Shilpi, and M Helal Uddin.** 2020. “Credit Rationing and Pass-Through in Supply Chains: Theory and Evidence from Bangladesh.” Technical report.
- Fabbri, Daniela, and Leora F Klapper.** 2016. “Bargaining Power and Trade Credit.” *Journal of Corporate Finance*, 41 66–80.
- Faccio, Mara.** 2006. “Politically connected firms.” *American Economic Review*, 96(1): 369–386.
- Fackler, Paul L, and Barry K Goodwin.** 2001. “Spatial price analysis.” *Handbook of Agricultural Economics*, 1 971–1024.
- Fafchamps, Marcel.** 2003. *Market institutions in sub-Saharan Africa: Theory and evidence.*: MIT press.
- Fafchamps, Marcel, Eleni Gabre-Madhin, and Bart Minten.** 2005. “Increasing returns and market efficiency in agricultural trade.” *Journal of Development Economics*, 78(2): 406–442.

- Fafchamps, Marcel, and Ruth Vargas Hill.** 2008. “Price transmission and trader entry in domestic commodity markets.” *Economic Development and Cultural Change*, 56(4): 729–766.
- Filmer, Deon, Jed Friedman, Eeshani Kandpal, and Junko Onishi.** 2018. *General equilibrium effects of targeted cash transfers: Nutrition impacts on non-beneficiary children.*: The World Bank.
- Fisman, Raymond, and Inessa Love.** 2003. “Trade credit, financial intermediary development, and industry growth.” *Journal of Finance*, 58(1): 353–374.
- Fisman, Raymond, and Mayank Raturi.** 2004. “Does competition encourage credit provision? Evidence from African trade credit relationships.” *Review of Economics and Statistics*, 86(1): 345–352.
- Fold, Niels.** 2005. “Global cocoa sourcing patterns.” *Cross-continental Agro-food Chains: Structures, Actors and Dynamics in the Global Food System*, p. 223.
- Gayi, Samuel K., and Komi Tsowou.** 2015. *Cocoa industry: Integrating small farmers into the global value chain.*: UNCTAD.
- Ghani, Tarek, and Tristan Reed.** 2021. “Relationships on the rocks: Contract evolution in a market for ice.” *American Economic Journal: Microeconomics*.
- Graddy, Kathryn.** 1995. “Testing for imperfect competition at the Fulton fish market.” *RAND Journal of Economics* 75–92.
- Hildebrandt, Nicole, Yaw Nyarko, Giorgia Romagnoli, and Emilia Soldani.** 2015. “Price Information, Inter-Village Networks, and “Bargaining Spillovers”: Experimental Evidence from Ghana.” Technical report, Working paper.
- Iacovone, Leo, and David McKenzie.** 2019. “Experimental Evidence on Shortening Supply Chains for Fruit and Vegetable Vendors in Bogota’s Slums.” Technical report, Working Paper.
- ICCO.** 2019. *Quarterly Bulletin of Cocoa Statistics.*, XLV(1): .
- IFAD.** 2018. *Project Completion Report Validation - Rural Finance and Community Improvement Programme, Republic of Sierra Leone.*
- Jaffe, Sonia, and E Glen Weyl.** 2010. “Linear demand systems are inconsistent with discrete choice.” *BE Journal of Theoretical Economics*, 10(1): .
- Jensen, Robert, and Nolan H Miller.** 2018. “Market Integration, Demand, and the Growth of Firms: Evidence from a Natural Experiment in India.” *American Economic Review*, 108(12): 3583–3625.
- Kalouptsidi, Myrto.** 2018. “Detection and impact of industrial subsidies: The case of Chinese shipbuilding.” *The Review of Economic Studies*, 85(2): 1111–1158.
- Khwaja, Asim Ijaz, and Atif Mian.** 2005. “Do lenders favor politically connected firms? Rent provision in an emerging financial market.” *Quarterly Journal of Economics*, 120(4): 1371–1411.

- Kreps, David M, and Jose A Scheinkman.** 1983. “Quantity precommitment and Bertrand competition yield Cournot outcomes.” *Bell Journal of Economics* 326–337.
- Lalive, Rafael, Camille Landais, and Josef Zweimüller.** 2015. “Market externalities of large unemployment insurance extension programs.” *American Economic Review*, 105(12): 3564–3596.
- Macchiavello, Rocco, and Ameet Morjaria.** 2015. “The value of relationships: evidence from a supply shock to Kenyan rose exports.” *American Economic Review*, 105(9): 2911–45.
- Macchiavello, Rocco, and Ameet Morjaria.** 2020. “Competition and Relational Contracts in the Rwanda Coffee Chain.” *Quarterly Journal of Economics*.
- Mach, Traci L, John D Wolken, Courtney M Carter, John A Holmes, and Lieu N Hazelwood.** 2006. “Financial services used by small businesses: Evidence from the 2003 Survey of Small Business Finances..” *J Federal Reserve Bulletin*, p. p.A167.
- Maitra, Pushkar, Sandip Mitra, Dilip Mookherjee, Alberto Motta, and Sujata Visaria.** 2017. “Financing smallholder agriculture: An experiment with agent-intermediated microloans in india.” *Journal of Development Economics*, 127 306–337.
- McMillan, John, and Christopher Woodruff.** 1999. “Interfirm relationships and informal credit in Vietnam.” *Quarterly Journal of Economics* 1285–1320.
- Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria.** 2018. “Asymmetric information and middleman margins: An experiment with Indian potato farmers.” *Review of Economics and Statistics*, 100(1): 1–13.
- Mobarak, Ahmed Mushfiq, and Mark Rosenzweig.** 2014. “Risk, insurance and wages in general equilibrium.” Technical report, National Bureau of Economic Research.
- Muralidharan, Karthik, and Paul Niehaus.** 2017. “Experimentation at scale.” *Journal of Economic Perspectives*, 31(4): 103–24.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar.** 2016. “General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence from India.” *Department of Economics, University of California, San Diego, processed.*
- Osborne, Theresa.** 2005. “Imperfect competition in agricultural markets: evidence from Ethiopia.” *Journal of Development Economics*, 76(2): 405–428.
- Ottaviano, Gianmarco, Takatoshi Tabuchi, and Jacques-François Thisse.** 2002. “Agglomeration and trade revisited.” *International Economic Review* 409–435.
- Panagariya, Arvind.** 2000. *Evaluating the case for export subsidies.*: The World Bank.
- Petersen, M.A., and R.G. Rajan.** 1997. “Trade credit: theories and evidence..” *The Review of Financial Studies*, 10(3): 661–691.
- Rijkers, Bob, Leila Baghdadi, and Gael Raballand.** 2015. “Political connections and tariff evasion evidence from Tunisia.” *The World Bank Economic Review*, 31(2): 459–482.

- Rodrik, Dani.** 1993. "Taking trade policy seriously: export subsidization as a case study in policy effectiveness." *NBER Working Paper*(w4567): .
- Rotemberg, Martin.** 2019. "Equilibrium effects of firm subsidies." *American Economic Review*, 109(10): 3475–3513.
- Rubin, Donald B.** 1974. "Estimating causal effects of treatments in randomized and nonrandomized studies.." *Journal of Educational Psychology*, 66(5): , p. 688.
- Sandefur, Justin, and Bilal Siddiqi.** 2013. "Delivering Justice to the Poor: Theory and Experimental Evidence from Liberia." 20.
- Sitko, Nicholas J, and TS Jayne.** 2014. "Exploitative briefcase businessmen, parasites, and other myths and legends: assembly traders and the performance of maize markets in eastern and southern Africa." *World Development*, 54 56–67.
- United Nations.** 2020. "Comtrade."
- Vives, Xavier.** 2001. *Oligopoly pricing: old ideas and new tools.*: MIT press.
- Weyl, E Glen, and Michal Fabinger.** 2013. "Pass-through as an economic tool: Principles of incidence under imperfect competition." *Journal of Political Economy*, 121(3): 528–583.

Figures

Figure 1: Map of study villages

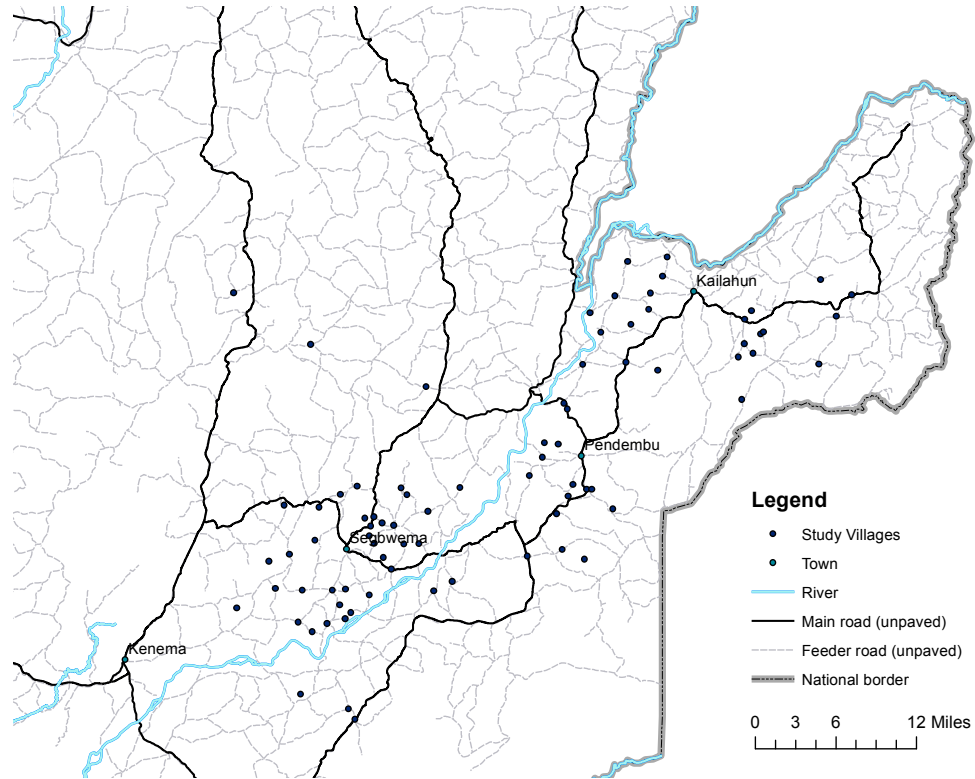
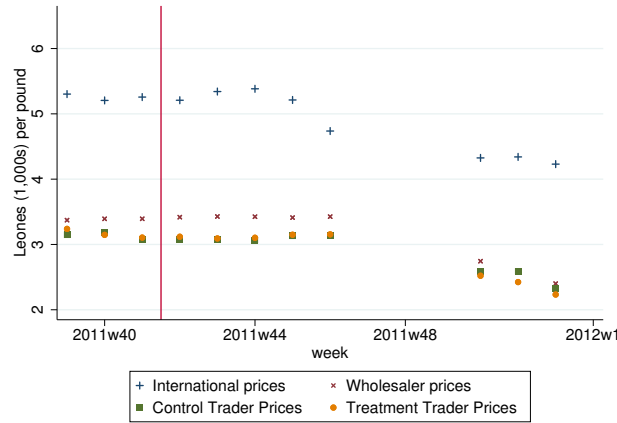
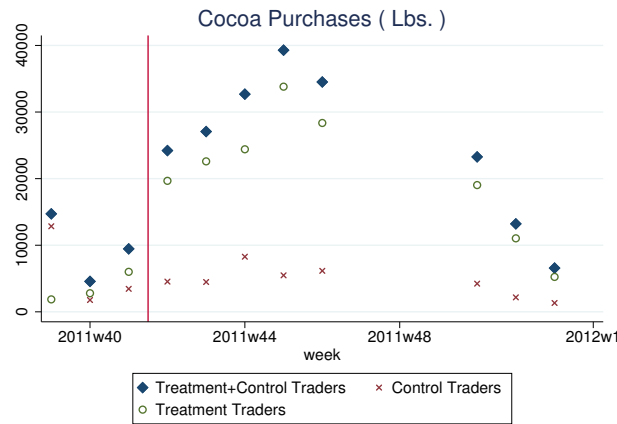


Figure 2: Cocoa Prices




Notes: The figure presents average weekly prices for: *i*) international cocoa prices; *ii*) prices the study wholesalers pay to the traders; *iii*) prices control traders pay to farmers; *iv*) prices treatment traders pay to farmers. All prices are in Leones (1,000s) per pound. Wholesale and farmer price data collection was suspended for most of three weeks (w47-w49). The vertical line marks the beginning of the intervention period.

Figure 3: Purchases of Cocoa



Notes: The figure shows the total amount of cocoa purchases by study traders (i.e., control and treatment traders), control traders, and treatment traders. The vertical line marks the beginning of the intervention period. Data collection was suspended for most of three weeks (w47-w49).

Figure 4: Counterfactual Experiments: Return on Investment by Γ



Notes: The graphs show the return on investment (ROI) in the three counterfactuals described in Section 5.5. The ROI is defined as the ratio between the additional farmer revenues induced by the intervention and the cost of the intervention. Panel A and B focus on the case of homogeneous traders. Panel A shows the ROI where 20% of traders are subsidized ($\mu = .2$). Panel B shows the ratio of the ROI with $\mu = .8$ to $\mu = .2$. In Panel C, half of the traders have a 5% higher resale price and are therefore “large”. The figure compares the ROI when half of both types of traders get the subsidy ($\mu_L = \mu_S = .5$) and when only large traders get the subsidy ($\mu_L = 1, \mu_S = 0$). In the three panels, these outcomes are plotted as a function of Γ , while keeping the pass-through rate constant at 0.91 and adjusting the number of firms, n , according to Equation (9).

Tables

Table 1: Baseline Trader Summary Statistics

Covariate	Treatment	Control	Treatment - Control
<i>Panel A: Baseline Interview</i>			
Self-estimate bags sold in 2011	20.0 (28.3)	18.6 (18.5)	1.5 (2.23)
Age, years	38.2 (8.2)	36.9 (10.2)	1.4 (1.91)
Years trading cocoa	8.1 (5.4)	8.9 (5.5)	-0.8 (1.2)
Years selling to study wholesaler	5.7 (4.8)	7.3 (4.9)	-1.6 (0.86)*
Cement or tile floor in house $\in \{0, 1\}$	0.53 (0.51)	0.63 (0.49)	-0.1 (0.1)
Mobile phone owner $\in \{0, 1\}$	0.90 (0.30)	0.93 (0.27)	-0.03 (0.06)
Access to storage facility $\in \{0, 1\}$	0.88 (0.33)	0.78 (0.42)	0.10 (0.09)
Villages operating in	4.25 (1.64)	4.87 (2.02)	-0.62 (0.39)
Number of suppliers per village	5.8 (3.3)	6.2 (3.6)	-0.35 (0.84)
Share of suppliers given credit since March	0.72 (0.32)	0.68 (0.28)	0.04 (0.05)
<i>Panel B: Pre-treatment shipment data</i>			
Price Paid to Farmer (shipment-level)	3,137 (154)	3,136 (151)	1.2 (41.9)
Pounds sold during pre-treatment (weekly)	89 (235)	151 (583)	-62 (41)

Notes: Panel A presents balancing for the variables defined in the baseline survey. Some baseline survey variables are missing for one trader. The column “Treatment-Control” presents results from a regression on treatment and randomization pairs. Panel B presents balancing for variables from pre-experiment shipment data. Prices are defined only for the subset of traders that delivers at least one shipment during this period. Quantities are defined for all traders and are equal to zero for traders who do not make any delivery in the pre-experimental period. Standard errors are clustered by trader. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2: Treatment-Control Differences in Farmer Prices

	(1)	(2)	(3)	(4)	(5)
Treatment Trader	-32.52	-5.47	-5.92	-12.87	-6.86
	(48.40)	(15.39)	(17.58)	(13.66)	(16.02)
Control Group Mean	2987	2987	2987	2987	2987
Week FE		X	X	X	X
Trader Controls			X		X
Village Controls				X	X
Observations	1079	1079	1060	1079	1060

Notes: The table reports the difference between the prices paid by treatment and control traders to farmers during the experiment, measured in Leones per pound. The subsidy to treatment traders was Le. 150. per pound. An observation is a shipment delivered by the trader to a wholesaler. Trader controls are baseline values of pounds of cocoa sold, number of villages operating in, number of suppliers buying from, share of clients given credit in baseline, age, years of working with wholesaler, and dummies for ownership of a cement or tile floor, mobile phone and access to a storage facility. Village controls are baseline share of suppliers begin given credit, number of other bonus traders and number of study traders, miles to nearest town, and number of clients across all traders. Data on some trader controls are missing for one trader and thus the number of observations falls in Columns (3) and (5). Standard errors are clustered at the level of the trader. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Treatment-Control Differences in Advance Payments

	(1)	(2)	(3)	(4)
Treatment Trader	0.12*** (0.03)	0.12*** (0.03)	0.12*** (0.03)	0.12*** (0.03)
Control Group Mean	0.15	0.15	0.15	0.15
Trader Controls		X		X
Village Controls			X	X
Observations	1818	1806	1818	1806

Notes: The table reports the difference between treatment and control in the share of regular suppliers that receive advance payments (binary indicator) during the experimental period. An observation is a farmer a trader listed as regular supplier in the baseline survey. Trader controls are baseline values of pounds of cocoa sold, number of villages operating in, number of suppliers buying from, share of clients given credit in baseline, age, years of working with wholesaler, and dummies for ownership of a cement or tile floor, mobile phone and access to a storage facility. Village controls are baseline share of suppliers receiving credit, number of other bonus traders and number of study traders, miles to nearest town, and number of clients across all traders. Data on some trader controls are missing for one trader and thus the number of observations falls in Columns (2) and (4). Standard errors are clustered at the level of the trader. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Treatment-Control Differences in Quantities

	(1)	(2)
Treatment Trader	398.38***	395.57***
	(38.0)	(38.3)
Control Group Mean	114.6	114.6
Trader Controls		X
Observations	640	632

Notes: The table reports the difference between the quantities of cocoa purchased by treatment and control traders during the experimental period. An observation is a week*trader (8*80). Trader controls are baseline values of pounds of cocoa sold, number of villages operating in, number of suppliers buying from, share of suppliers receiving credit from the trader at baseline, age, years of working with wholesaler, and dummies for ownership of a cement or tile floor, mobile phone and access to a storage facility.

Data on some trader controls are missing for one trader and thus the number of observations falls in Column (2). Standard errors are clustered at the level of the trader. ***p<0.01, **p<0.05, *p<0.1.

Table 5: Pass-Through from Wholesale Price to Farmer Prices

	(1)	(2)	(3)
Wholesale Price	0.91***	0.91***	0.85***
	(0.06)	(0.06)	(0.03)
Control Group Mean	3007	3007	2960
Kleibergen-Paap First Stage F-stat	57.3	62.0	32.5
Trader FE	X	X	
Village FE		X	
Observations	1254	1254	72

Notes: The table reports the pass-through from wholesaler prices to farmer prices. Both are measured in Leones per pound. In all columns, wholesaler prices are instrumented with the front-month prices for liquid cocoa futures, obtained by the *Intercontinental Exchange*. In Columns (1)-(2), an observation is a shipment delivered by the trader to a wholesaler before or during the intervention and standard errors are clustered by trader and date. In Column (3), we collapse data in a time-series of average prices for each date and use heteroskedasticity and autocorrelation consistent (HAC) standard errors, with Newey-West kernels. ***p<0.01, **p<0.05, *p<0.1.

Table 6: Estimates of the Model Parameters

	(1)	(2)	(3)	(4)	(5)	(6)
	λ	δ	$p_T - p_C$	Γ	β	n
Panel A: Main Results						
1. <i>Cross-Sectional Variation in Prices and Credit</i>	.95	.95	12.8 [-17.8,47.9]	.09 [-.11,.47]	.35 [.24,.48]	11.1 [8.8,18.9]
Panel B: Sensitivity						
2. <i>Treatment Heterogeneity in Prices and Credit</i>	.93	.93	20.3 [-10.5,55.6]	.16 [-.07,.59]	.32 [.22,.45]	12 [9.1,23.4]
3. <i>Calibration of Interest Rate</i>	.98	.98	1.7 [-28.6,36.7]	.01 [-.16,.32]	.37 [.26,.51]	10.3 [8.5,15.5]
4. <i>Farmer Time-Preference Experiment</i>	.95	.914	27.2 [-3.75,62.6]	.22 [-.12,.72]	.31 [.21,.43]	12.8 [9.3,30.2]
5. <i>Only Grade-A Credit</i>	.92	.92	19.6 [-11,54.7]	.15 [-.07,.57]	.33 [.23,.45]	11.8 [9.1,22.7]

Notes: The table presents results from the estimation in Section 5 (further described in Appendix B). The rows differ in the strategy to calibrate λ and δ . Panel A presents results of our baseline approach, in which we derive λ based on baseline cross-village variation in prices and credit and assume $\lambda = \delta$. Panel B presents results from four alternative approaches to compute λ and δ . For the parameters estimated in Columns (3)-(6), we report in square brackets 90% confidence intervals from bootstrapping at the randomization pair level.

Appendix [FOR ONLINE PUBLICATION]

A Measuring Cocoa Quality

Both international and local cocoa prices vary with quality. Factors contributing to poor quality cocoa are high moisture content, mold, germination, a lack of fermentation and slate, a discoloration signaling poor flavor. There is wide agreement on these standards internationally. For a discussion, refer to CAOBISCO/ECA/FCC (2015) and, for the specific case of West Africa, David (2005). Other dimensions of quality affecting price on the international market are various fair-trade and environmental certifications. Such certification generally requires that beans can be verifiably traced to individual producers. In our market, there is not yet the infrastructure to do such tracing, and so this quality dimension does not apply.

In our grading system, inspectors from our research team with local language skills stayed in the warehouses of wholesalers and tested a sample of 50 beans from each bag of cocoa as it arrived. Moisture was measured using Dickey John MiniGAC moisture meters, two of which were generously donated by the manufacturer. Other defects were spotted by eye, after cracking beans open with a knife. Grade A beans have no more than average 11.5% moisture, no more than 2% mold (1 bean of 50), and no less than 72% beans with no defect (36 beans of 50). Grade B beans have no more than 22% moisture, 4% mold (2 beans of 50) and no less than 52% good beans (27 beans of 50). Grade C applies to any bean failing to be grade A or B.

B Computing Effective Prices

B.1 Defining the Effective Price

A farmer can sell cocoa at price p_1 at harvest time or at $p_0 = \lambda p_1$ in advance, where λ depends on the market interest rate, r : $\lambda = \frac{1}{1+r} < 1$. The farmer will prefer the advance if $\frac{\lambda}{\delta} p_1 > p_1$, where δ is the farmer's discount factor between the two periods. Intuitively, if the discount factor is low (i.e. the farmer is impatient) or the penalty for getting early payment is low (i.e. λ is high), then the farmer prefers the advance.

A trader who pays a share s of the cocoa in advance, will thus pay the average price:

$$\tilde{p} = sp_0 + (1-s)p_1 = p_1(s\lambda + (1-s)) = p_1(1 - s(1 - \lambda)).$$

We then denote with p the *effective price* paid by each trader, *as valued by the farmer* (at harvest time):

$$p = s\frac{p_0}{\delta} + (1-s)p_1 = p_1(s\frac{\lambda}{\delta} + (1-s)) = p_1(1 - s(1 - \frac{\lambda}{\delta})) = \frac{\tilde{p}}{1 - s(1 - \lambda)}(1 - s(1 - \frac{\lambda}{\delta}))$$

The effective price p shapes farmers' supply and it is our object of interest.

B.2 Computing the Effective Price

In the data, we observe the average price paid by traders in each treatment group, \tilde{p} . We proxy s with the share of farmers who receives credit. We compute the effective price (by treatment group) following several approaches. In Section 5.1.1, we described our baseline approach based on cross-sectional variation in prices and credit. Here, we present details on our alternative strategies.

Treatment Heterogeneity in prices and credit. We infer the value of advance payments from the covariance of treatment-control differences in prices and treatment-control differences in advance payments. Similarly to our baseline approach, we assume again that $\lambda = \delta$. The slope between the price and credit responses thus identifies their relative value, or how much less a trader who increases his advance payments needs to adjust his prices.

For this purpose, we modify Equation (11) to allow for heterogeneity in the treatment-control

differences across villages and trader characteristics:

$$AdvancePayment_{fizv} = \eta_z + \pi^a(\text{Treat}_i) + (\text{Treat}_i \times X'_v)\pi_v^a + X'_v\beta_v + (\text{Treat}_i \times X'_i)\pi_i^a + X'_i\beta_i + \nu_{sipv}, \quad (\text{B.1})$$

where X_v is the vector of village covariates and X_i is a vector of trader covariates. For any trader-village pair iv we then compute the predicted treatment-control difference in advance payment provision using heterogeneity by X_v and X_i : $\widehat{DTC}_{iv}^a = X'_v\pi_v^a + X'_i\pi_i^a + \pi^a$. Finally, we run the following specification to test whether village-trader pairs with larger treatment-control differences in advance payments display lower differences in prices:

$$\tilde{p}_{sizvt} = \eta_z + \eta_t + \pi^{\tilde{p}}(\text{Treat}_i) + \pi_a^{\tilde{p}}(\widehat{DTC}_{iv}^a \cdot \text{Treat}_i) + X'_i\beta_i + X'_v\beta_v + \epsilon_{kiptv}. \quad (\text{B.2})$$

If total payments and advance payments are substitutes (i.e., $\lambda > 0$), then $\pi_a^{\tilde{p}} < 0$.²⁰

Appendix Figure B.1 provides presents initial evidence that there is a negative slope between the treatment-control differences along the two margins. Here we estimate treatment-control differences in prices and advance payments in each of the chiefdoms included in the study, and plot them against each other. Chiefdoms are geographic units of local legal and political administration, and, as discussed in Acemoglu et al. (2014) vary in contract enforcement and other institution (unfortunately, our data do not include explicit information on contract enforcement institutions). The scatter displays a negative relation: the regression line has a slope of -199.

Appendix Table B.2 presents estimates of $\pi_a^{\tilde{p}}$. In the different columns we show estimates generated using different sets of controls to predict \widehat{DTC}_{iv}^a . Since \widehat{DTC}_{iv}^a is an estimated regressor, we follow Bertrand et al. (2004) and Cameron et al. (2008) and present p-values calculated using the bootstrap-t procedure (Efron, 1981). We draw 2,000 bootstrap samples, clustering the bootstrapping by randomization pair. Our estimates of $\pi_a^{\tilde{p}}$ are negative and statistically significant at 7 to 15 percent across the three specifications. In column (1) of Table B.2, \widehat{DTC}_{iv}^a is predicted using only chiefdom dummies. The estimate using these dummies predicts that a village where treatment traders are 12 percentage points more likely to provide advance payments than control traders—the mean coefficient in Table 3—would have a treatment-control difference in prices that is 40 Leones lower than a village with no difference in advance payments. This is economically meaningful as it accounts for a reduction in the treatment difference of about one-quarter of the subsidy value. In Appendix Table B.2, we find similar results in column (2), where the effect on advance payments is predicted using chiefdom dummies and village covariates, and in column (3), where we also add

²⁰Since \widehat{DTC}_{iv}^a is collinear with the vector of controls, its level is not included in the estimating equation.

trader covariates. While the magnitude of the coefficients falls across columns, the core result holds: price and advance payment responses are substitutes. According to column (3), when treatment traders provide credit but control traders do not, shipment prices would fall by 221 Leones, from a baseline of 3,138 Leones. Thus, $\lambda = 0.93$.

Interest Rate Calibration According to the World Development Indicators, the average lending rate for Sierra Leone over the last fifteen years was between 21% and 25% per year. In the inventory credit evaluation described by Casaburi et al. (2014), rates on subsidized collateralized loans for agricultural smallholders were 22%. One would expect rates on unsecured agricultural loans to be higher (IFAD, 2018). As a lower bound, we consider a rate of 2% per month. We assume a loan duration of one month, which is reasonable because we focus only on credit provided during the intervention (2.5 months) and most transactions happen in the first half of the intervention. Thus, $\lambda = 0.98$ and the difference in effective prices between treatment and control is 0.01.

Experimental Elicitation of Farmers’ Subjective Rates of Returns. We conducted a lab-in-the-field experiment with cocoa farmers. Similarly to the original field experiment, the study took place in the harvest season (in November). It targeted three villages included in the original experiment. In each village, we listed approximately thirty cocoa farmers. Farmers were then asked to choose between a payment in three days and one in 24 or 45 days, i.e. three weeks or six weeks after the first payment).²¹ We follow a *Multiple Price List* design (Andersen et al., 2008). For each time interval, farmers made 15 binary choices between the earlier and the later payment, with the latter increasing across choices. The earlier amount was always Leones 300,000 (approximately USD 30). This is a large amount, worth about 7% of the median self-report cocoa revenues from the entire harvest season. Future amounts spanned between Leones 300,000 and 700,000.

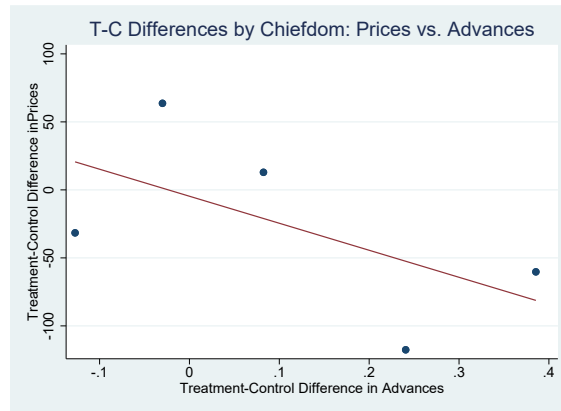
We then use the switching points in each set of choices to compute a farmer’s monthly *Required Rate of Return*, RRR . As highlighted in Cohen et al. (2020), the RRR is not necessarily a time-preference parameter because a number of other factors may affect its determination (e.g. risk aversion and liquidity constraints). Rather, it is simply an indifference point in the experimental data (we define the indifference point as the mid point between the highest future amount to which the subject prefers the (fixed) present amount and the lowest future amount the subject prefers to the present amount). Nevertheless, the RRR is still the relevant object to understand how much farmers may value traders’ advance payments. Pooling our data, we compute a median monthly RRR of 0.0935, which implies $\delta = 0.914$.

²¹To avoid differences in transaction costs and perceived risk, we add a front-end delay in the earlier. This design also advance provision, as traders often return with the cash few days after agreeing on the advance transaction. We also elicited time-preferences between 24 and 45 days.

Scaling Credit by the Share of Grade-A Transactions. In the final sensitivity test, we focus on the observation that some of the advance payments may concern non-grade A transactions (as per our discussion in Section 4.1.2). First, we first scale down the village-level credit share by the share of grade-A transactions in that village. Moving from a village where no farmer receives advance payments at baseline to a village where each farmer receives advance payments decreases shipment prices paid by the trader by approximately 254 Leones from an average of 3,138 , i.e. $\lambda = 0.92$. Intuitively, this value of λ is lower than when we use the unadjusted credit share. Second, we multiply the share of advance payments by the share of grade-A transactions in each treatment group, thus obtaining $\sigma_C = 0.15 * 0.45 = 0.068$, $\sigma_T = 0.27 * 0.63 = 0.17$. This gives a difference in effective prices between treatment and control traders equal to 19.6 Leones [-11,54.7].

B.3 Figures

Figure B.1: Estimating λ : Treatment-Control Differences by Chiefdom, Prices vs. Advance Payments



Notes: The scatter reports the correlation across price and advance payments treatment-control differences, estimated separately across the five chiefdoms included in the study. The regression line has a slope of -199.

B.4 Tables

Table B.1: The Value of Advance Payments: Baseline Correlations

	(1)	(2)
Share of Farmers Receiving Advance Payments	-149.65*	-147.19*
	(74.66)	(75.47)
Dependent Variable Mean	3138	3138
Village Controls		X
Observations	43	43

Notes: The table presents correlation between baseline value of the average village cocoa price and the share of farmers receiving advance payments in the village. The sample includes 44 villages for which we have baseline cocoa shipment data. Village controls include: baseline share of suppliers receiving credit, number of traders in the village, distance from the wholesaler warehouse, and number of farmers in the village. Standard errors allow for heteroskedasticity. ***p<0.01, **p<0.05, *p<0.1.

Table B.2: The Value of Advance Payments: Heterogeneity in Treatment-Control Differences

	(1)	(2)	(3)
Treat* Estimated Treatment Effect on Credit	-328	-273	-221
<i>p-values from bootstrapped t-stats</i>	[.07]	[.14]	[.15]
Chiefdoms	X	X	X
Village Controls		X	X
Trader Controls			X
Observations	1060	1060	1060

Notes: The dependent variable is the price paid by the trader for the shipment of cocoa. Each column presents estimates of $\pi_a^{\hat{p}}$ from equation B.2. P-values in brackets are derived from pairs cluster bootstrap-t at the randomization pair level using 2,000 replications. Trader controls are baseline values of pounds of cocoa sold, number of villages operating in, number of suppliers buying from, share of suppliers receiving credit from the trader at baseline, age, years of working with wholesaler, and dummies for ownership of a cement or tile floor, mobile phone and access to a storage facility. Village controls are number of other bonus traders and number of study traders, miles to nearest town, and number of clients across all traders.

C Alternative Estimation Moments

This Appendix presents details about the alternative approach to recover Γ and n . Our goal is to identify alternative moments and to compare the results we obtain from these moments to the ones of the main approach presented in the paper. Showing that different moments deliver similar estimates would provide support for the specific model we use.

C.1 Methodology

In the paper, we showed how they two parameters could be identified, relying on two moments: the *level* difference in treatment and control prices (Equation 5) and the pass-through rate of changes in wholesaler prices (Equation 9). In this section, we show how the key parameters Γ and n , and also the intercept parameter α , can be recovered from the *percent* differences between treatment and control in prices *and quantities*, combined again with the pass-through rate.

First, we derive theoretical expressions for the percent differences between treatment and control in prices and quantities:

$$\% \Delta p \equiv \frac{p^T - p^C}{p_C} = \frac{s\Gamma(1 + (1 - \Gamma)(n - 1))}{(1 - \Gamma)\mu ns + (1 + \Gamma)((1 - \Gamma)(n - 1)v + (\alpha + v))} \quad (\text{C.1})$$

and

$$\% \Delta q \equiv \frac{q_T - q_C}{q_C} = \frac{s(-2 - (1 - \Gamma)(n - 1))}{(1 - \Gamma)\mu ns - (1 + \Gamma)(v - \alpha)} \quad (\text{C.2})$$

For a given value of the subsidy s , these expressions depend on additional parameters, i.e., μ, v, α , as well as on those we aim to recover, i.e., Γ and n . We calibrate the value of μ and v . We set the former at 1/5, the share of treatment traders out of the total number of traders (study and non-study). Assigning a value to the latter requires some additional assumption. The (average) value of the wholesaler price (i.e. the price at which traders resell), is Le. 3,260. The average effective price at which traders purchase cocoa from farmers is Le. 3,010, 92% of the wholesaler price. However, in the model, v is the *net* resale price, net of other costs the traders may incur and that we do not observe, such as transport and storage costs. We set $v = 3,145$, which implies a 5% markdown.²²

C.2 Results

Having assigned values to μ and v , we have a system of three equations—Equations C.1, and C.2 defined above and the pass-through formula (Equation 9)—, in three unknowns, Γ, n , and α . We note that the intercept term α is identified only up to the currency unit choice.

During the experiment, control traders pay an average effective price of 3,010 Leones and, in our preferred specification with $\lambda = 0.95$, treatment traders pay 3,022.8. This implies that the percent price difference between treatment and control traders during the experiment is 0.4%. The average quantity purchased by control traders is 114.6 kilograms and the treatment-control difference is 395 kg. Thus, the percent difference between treatment and control traders is 344%.

Solving the equation system with these values for $\% \Delta p$ and $\% \Delta q$, we obtain the following estimates for the three parameters of interests: $\Gamma = 0.096$, $n = 11.08$, and $\alpha = 2,430$. The results for Γ and n are thus very close to the ones obtained when using the more parsimonious methodology described in the main text. We see this as evidence in support of the specific competition model chosen for the analysis.


Finally, we emphasize that the similarity of the results between the two approaches is not a mechanical result. This is because one set of results uses the *level* of the difference between treatment and control prices, while the other uses the *percent* differences between treatment and control in both prices and quantities. Figure C.1 and C.2 confirm

²²Results are quite stable when using other values of v , spanning between 3,010 (the average effective price paid to the farmer) and 3,260 (the average wholesaler price).

this point: the two graphs show, respectively, how the estimated values of Γ and n would vary with different values of the percent treatment-control difference in prices, $\frac{p_T - p_C}{p_C}$, in a neighborhood of the real value, 0.007 (represented by the vertical gray line). In each graph, the large dot reports the estimate from the main estimation presented in the text. The key point is that, while the estimates derived when using the real value $\frac{p_T - p_C}{p_C}$ are close to those in the main text, they would be quite different when using arbitrary values of $\frac{p_T - p_C}{p_C}$ (i.e. if the treatment-control difference in the level of prices were equivalent to a different value of the difference in percent terms.).


C.3 Figures

Figure C.1: Sensitivity of Γ to $\frac{p_T - p_C}{p_C}$



Notes: The graph reports sensitivity of the estimate of Γ obtained from the method described in Appendix C to the value of the percent treatment-control price difference. The dot represent the estimate from the main method presented in Section 5.

Figure C.2: Sensitivity of n to $\frac{p_T - p_C}{p_C}$



Notes: The graph reports sensitivity of the estimate of Γ obtained from the method described in Appendix C to the value of the percent treatment-control price difference. The dot represent the estimate from the main method presented in Section 5.2.

D The Impact of the Experiment on Prices, Quantities, and Farmer Revenues

This Appendix provides details on the steps to assess the impact of the experiment on prices and quantities with respect to a counterfactual without the experiment (Section 5.3.2).

Setup

We use the superscript 0 to refer to the scenario without the intervention and superscript 1 for the intervention. Without the experiment, traders are homogeneous and pay p^0 .²³ Each trader thus faces the direct supply $q_i^0 = a + bp_i^0 - c \sum_{i \neq j} p_j^0$.²⁴ Symmetry implies $q^0 = a + (b - c(n - 1))p^0$. Aggregate supply is thus $Q^0 = nq^0 = n[a + (b - c(n - 1))p^0]$. Throughout this section, we assume that non-study traders are equal to control ones.

Impact on Prices

To assess the impact of the experiment on the prices of control and treatment traders, we first compute the derivative of equilibrium prices with respect to the subsidy: $\frac{\partial p_T}{\partial s} = \frac{\Gamma - \frac{(\Gamma-1)\mu n}{\Gamma n + n + 1}}{\Gamma + 1}$, $\frac{\partial p_C}{\partial s} = \frac{(1-\Gamma)\mu n}{(\Gamma+1)(1+\Gamma+n(1-\Gamma))}$.

The impact of the experiment on prices is then given by $dp_g = \frac{\partial p_g}{\partial s} s$, for $g = \{T, C\}$. Given our estimates of Γ and n , we can compute the (rounded) $dp_T = p_T^1 - p^0 = 37.5$ and $dp_C = p_C^1 - p^0 = 24.7$.²⁵ Using a baseline price of $p^0 = 2,985$ (mean of the effective price for control traders during the experiment minus dp_C), we obtain $p_T^1/p^0 = 1.0125$ and $p_C^1/p^0 = 1.008$. The experimental subsidy, which was worth about 5% of the baseline price, increased treatment (control) prices by around 1.2% (0.8%). The change in aggregate quantity is given by $\mu dq_T + (1 - \mu) dq_C$, which corresponds to 4.1% of the baseline quantity.

Impact on Quantities

Given dp_T and dp_C , we can write: $q_T^1 = a + b(p^0 + dp_T) - c((\mu n - 1)(p^0 + dp_T) + (1 - \mu)n(p^0 + dp_C))$. The derivative of equilibrium quantities with respect to the subsidy are therefore: $\frac{\partial q_T}{\partial s} = -\frac{(1-\Gamma)\mu n}{\beta(1+\Gamma)(1+\Gamma+(1-\Gamma)n)}$ and $\frac{\partial q_C}{\partial s} = \frac{\Gamma+(1-\Gamma)(1-\mu)n+1}{\beta(1+\Gamma)(1+\Gamma+(1-\Gamma)n)}$.

The impact of the experiment on prices is then given by $dq_g = \frac{\partial p_g}{\partial s} s$, for $g = \{T, C\}$. Given our estimates of Γ and n , we can compute the (rounded) $dq_T = q_T^1 - q^0 = 324.2$ and $dq_C = q_C^1 - q^0 = -71.3$. Using a baseline quantity of $q^0 = 185.9$ (average quantity for control traders during the experiment minus dq_C), we obtain $q_T^1/q^0 = 2.74$ and $q_C^1/q^0 = 0.62$.

Return on Investment

We consider the return on investment (ROI) on interventions that treat a share μ of traders. We focus on a social planner whose welfare is linear in farmer revenues (and does not depend on trader revenues). Therefore, the ROI is the ratio between the increase in farmer revenues and the cost of the program.

In the pre-experiment period, farmer revenues are: $r^0 = p^0 Q^0 = p^0 n q^0$. In the experimental period, these become $r_1 = n(\mu p_T^1 q_T^1 + (1 - \mu) p_C^1 q_C^1)$. The cost of the intervention is $C = s n \mu q_T^1$. The ROI is thus $\frac{r_1 - r^0}{C}$.

²³Thus, throughout the exercise, we assume non-study traders and study traders are homogeneous before the experiment.

²⁴The direct supply function is $q_i = a + bp_i - c \sum_{j \neq i} p_j$, with $a \equiv \frac{\alpha}{\beta + \gamma(n-1)}$, $b \equiv \frac{\beta + \gamma(n-2)}{(\beta + \gamma(n-1))(\beta - \gamma)}$, $c \equiv \frac{\gamma}{(\beta + \gamma(n-1))(\beta - \gamma)}$.

²⁵This is, by construction, consistent with our estimate of the difference in (effective) prices between treatment and control traders.

E Model Extensions

E.1 Trader Heterogeneity

The model presented in Section 2 assumes that traders are symmetric at baseline and that the experimental subsidy is the only source of heterogeneity. The key results of the model, and thus the empirical strategy to recover the competition parameters, are robust to extensions that account for different forms of heterogeneity.

First, we allow baseline differences across traders in their resale prices.²⁶ For simplicity, we consider a case with two types of traders. Absent the experiment, a share σ of traders has resale price v , and a share $1 - \sigma$ has resale price $v' = v + w$. With the experiment, a share μ of traders in each group receives a per unit subsidy s . In equilibrium, firms with higher resale prices purchase larger quantities and pay higher prices (unless $\Gamma = 0$). By randomization, treatment is uncorrelated with firm characteristics. This orthogonality is the key benefit of randomization even if, as we discuss in the paper, the SUTVA is violated.

Within each group of traders (v and v'), the difference in equilibrium prices between treatment (subsidized) and control (unsubsidized) firms is $\Delta p = \frac{s\Gamma}{1+\Gamma}$. Therefore, trivially, this is the value for the expected price difference: $E[\Delta p] \equiv E[p_T - p_C] = \frac{s\Gamma}{1+\Gamma}$. Similarly, it can be shown that $E[\Delta q] \equiv E[q_T - q_C] = \frac{s}{\beta(1+\Gamma)}$. Finally, the linear inverse supply implies constant pass-through: For each type of firm, $\rho \equiv \frac{\partial p}{\partial v} = 1 - \frac{1}{1+\Gamma+n(1-\Gamma)}$, and thus $E[\rho]$ takes the same value. Therefore, the key moments presented in Equations (5) and (9) are unchanged.

Second, we allow for multiple differentiation rates across traders. We consider again a simple case with two groups of competitors. In a symmetric environment with n traders, each trader has $\frac{n}{2} - 1$ “close” competitors, with substitution rate γ , and $\frac{n}{2}$ “far” competitors with substitution rate $\kappa\gamma, 0 < \kappa < 1$. Therefore, the inverse supply for each trader i is $p_i = \alpha + \beta q_i + \gamma(\sum_{j \in C} p_j + \kappa \sum_{j \in F} p_j)$, where C and F represent close and far competitors, respectively.

It can be shown that the equilibrium differences between treatment and control are unchanged: $\Delta p = \frac{s\Gamma}{1+\Gamma}$ (where Γ is still $1 - \frac{\gamma}{\beta}$ and $\Delta q = \frac{s}{\beta(1+\tilde{n})}$). In addition, the pass-through rate is $\rho = 1 - \frac{1}{1+\Gamma+\tilde{n}(1-\Gamma)}$, where $\tilde{n} \equiv \frac{n}{2}(1 + \kappa)$ can be again defined as the “effective market size”, the number of competitors weighted by their (relative) substitution parameter κ . In this case, the estimation procedure presented in the paper therefore recovers Γ and \tilde{n} .²⁷

E.2 Non-study Traders

As discussed above, the model presented in Section 2 features symmetric traders. From this pool of identical traders, a share μ receives the experimental subsidy. In our field experiment setting, about 60% of the traders are not included in the study (and we do not collect data on them). These traders may be fundamentally different than the ones we include in the study. We present an extension of the model that accounts for this issue.

There is a share σ of study traders (S) and a share $1 - \sigma$ of non-study traders (NS). We allow the two types of farmers to vary in their resale prices: $v_S = v$ and $v_{NS} = v + w, w \neq 0$. Inverse supply for trader i is again $p_i = \alpha + \beta q_i + \gamma \sum_{j \neq i} q_j$.²⁸ A share μ of the study traders, and thus a share $\mu\sigma$ of all traders, receives the subsidy.

Our experimental estimates only compare prices of the study traders. The main object of interest is $p_{ST} - p_{SC}$, where the subscript S refers to the share σ of study traders. The moments derived in Section 2 are robust to the presence of non-study traders. It can be shown that $\Delta p_S \equiv p_{ST} - p_{SC} = \frac{s\Gamma}{1+\Gamma}$. This is the same value we obtained in the baseline model, where we assumed that all traders were part of the experiment (Equation 5). A similar result is obtained for Δq_S . Finally, the pass-through rate is also unchanged (again, this is due to the common pass-through functional form).

²⁶This is equivalent to varying producer costs in an oligopoly model.

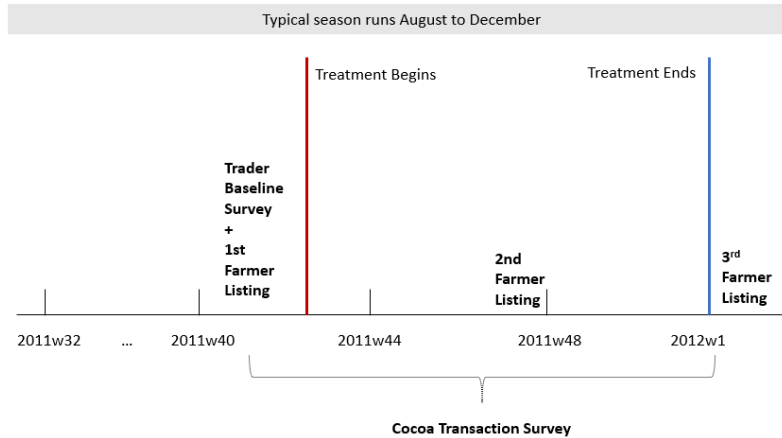
²⁷The result extends to the general case of $m = 1, \dots, M$ groups of traders, with differentiation $\Gamma_m = \kappa^{m-1}\Gamma$. In this case, Δp , Δq , and ρ are as above and $\tilde{n} = \frac{n}{M} \frac{1-\kappa^M}{1-\kappa}$.

²⁸That is, we assume a common degree of differentiation across study and non-study traders.

F Additional Figures and Tables

F.1 Timeline

Figure F.1: Timing of Experiment and Fielding of Survey Instruments



Notes: The figure shows the timeline of the intervention, the baseline survey, and the farmer listings. Traders were informed of the subsidy when the treatment began. Credit is offered and prices are negotiated throughout the crop harvest, and so we cannot indicate these moments on the timeline with accuracy. Our analysis of credit focuses only on credit offered during the 2.5 months of the experiment.

F.2 Price Results for Grade B and C

Table F.1: Treatment-Control Differences in Prices (Grade B and C)

	Grade B	Grade C	Grade B&C
	(1)	(2)	(3)
Treatment Trader	40.07*** (13.22)	-0.36 (18.24)	31.06** (13.72)
Control Group Mean	3025.82	3050.47	3035.91
Observations	532	231	763

Notes: The dependent variables in Columns 1, 2, 3 are the price per pound for grade B, grade C, and grade B&C, cocoa respectively. The subsidy to treatment traders for grade A was Le. 150. per pound. An observation is a shipment delivered by the trader to a wholesaler. All the regressions include week fixed effects. Standard errors are clustered at the level of the trader. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

F.3 Traders' Costs


Table F.2: Treatment-Control Differences in Costs

	Transport Costs	Truck Dummy
	(1)	(2)
Treatment Trader	-13.88*** (4.02)	0.19*** (0.06)
Control Group Mean	49.12	0.40
Observations	1079	1079

Notes: In Column (1), the dependent variable is the transport cost per pound, defined as the ratio of the transport cost per bag reported by the trader and the weight of the bag (measured in Leones per pound). In Column (2), the dependent variable is a dummy equal to one if the trader reports using a truck to transport that bag. The subsidy to treatment traders was Le. 150. per pound. An observation is a shipment delivered by the trader to a wholesaler. All the regressions include week fixed effects. Standard errors are clustered at the level of the trader. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

F.4 Price responses to subsidies to different share of traders

Figure F.2: Price responses to subsidies to different share of traders



Notes: The graph shows the impact of counterfactual experiments on (effective) prices paid by control traders (continuous line) and treatment traders (dashed line). Specifically, it reports the increase in prices in response to a unit-subsidy as a function of the share of treated traders, μ . The vertical line reports the share of traders treated in our experiment, $\mu = 0.2$. For $\mu \rightarrow 1$, the response of treatment traders tends to the pass-through rate, 0.91.