

Competition in Agricultural Markets: An Experimental Approach*

Lorenzo Casaburi Tristan Reed

October 2017

Abstract

This paper develops an experimental approach to measure competition among intermediaries in agricultural markets, based on the random allocation of subsidies to traders. We show that, in individual-level randomizations with competitive spillovers, treatment-control differences in prices can inform an intuitive test of the degree of differentiation among firms. In the context of the Sierra Leone cocoa industry, traders compete by providing farmers credit, as well as through prices. Even when accounting for both the price and the credit margin, differentiation among traders is low. By combining the experimental results with quasi-experimental estimates of the pass-through rate, we then estimate market size—the effective number of traders competing for farmers’ supply—and we find it to be substantially larger than the village. These results are consistent with a view of competitive agricultural markets.

*Lorenzo Casaburi: lorenzo.casaburi@econ.uzh.ch. Tristan Reed: reed.tristan@gmail.com. Previous versions of the paper were circulated with the title “Interlinked Transactions and Pass-Through: Experimental Evidence from Sierra Leone.” We thank Philippe Aghion, Pol Antràs, David Atkin, Dave Donaldson, Pascale 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 Center and the Agricultural Technology Adoption Initiative, and the institutional support of Innovations for Poverty Action in Freetown.

1 Introduction

The degree to which intermediaries compete is a long-standing object of interest in studies of agricultural markets in developing countries. Competition shapes how price signals propagate along supply chains, and the welfare implications of taxes and subsidies for producers and consumers. In this paper, we propose an experimental approach to estimate the degree of competition. Our experiment, which takes place in the Sierra Leone cocoa industry during the 2011 harvest season, is based on the randomization of unit subsidies to competing traders for their purchases from farmers. The main idea is that the difference in the prices that subsidized and unsubsidized firms pay is informative of the degree of differentiation among them. The basic estimation procedure uses only price data collected over one season, and thus presents advantages relative to other methods that require detailed trader cost data, a large cross section of markets, or a long time series.

In two influential papers, Weyl and Fabinger (2013) and Atkin and Donaldson (2015) show that, in a general model of symmetric imperfect competition, the pass-through rate—i.e., the difference in prices between two markets with an infinitesimal difference in costs for all producers—is the key object to recover the degree of competition in the market. Experimental or quasi-experimental analysis of pass-through thus requires studying plausibly independent markets that experience different cost shocks. However, researchers face common constraints in the design of market-level experiments. Defining truly independent markets can easily imply that the cluster size is so large that there are too few clusters (or, in the extreme case, a single cluster containing all firms), or that randomization across markets would be very expensive. This paper shows how an experimental approach can be valuable even in these cases. Using solely individual-level randomization (i.e., without variation in the cluster treatment intensity), it provides an intuitive interpretation of the treatment-control difference between subsidized and unsubsidized firms *when there are competitive spillovers across the two groups of firms*.

A simple framework of oligopsonistic competition among (potentially) differentiated traders guides the empirical analysis. In the framework, farmers sell their output to traders (at the “trader’s price”) and traders resell it to wholesalers (at the “wholesaler’s price”). The differentiation parameter captures a trader’s ability to buy at a lower price than her competitors.¹

¹In the model, the differentiation rate is related to the extent to which prices paid by competitors affect the quantity supplied to a trader, relative to the trader’s own price sensitivity. More precisely, since our framework features quantity-setting traders, the differentiation rate depends on the ratio between the slope

The model highlights the strategic interaction between subsidized and unsubsidized traders that compete with each other: in equilibrium, unsubsidized traders, as well as subsidized ones, adjust their behavior in response to the subsidy. Therefore, differences between treatment and control during the experiment cannot be interpreted as Rubin (1974) treatment effects. However, randomization ensures that subsidies are uncorrelated with firm characteristics and thus that the treatment-control price differences arise only from the experimental subsidy.

In the first contribution of the paper, we thus propose a simple relationship between the treatment-control difference in the prices that traders pay to farmers during the experiment and the differentiation rate across traders. Intuitively, there can be no systematic difference between the prices paid by subsidized and unsubsidized traders if traders are perfectly homogeneous from the farmers’ perspective. In this case, there will be only one “market price”. If traders are differentiated, however, different prices can coexist and unsubsidized traders can pay farmers a lower price than subsidized ones do. Through this insight, our experimental results can recover the degree of differentiation among traders.

Differentiation among traders can emerge for a several reasons. Recent literature has emphasized the role of search costs (Jensen, 2007; Aker, 2010; Allen, 2014, Startz, 2016). In addition, lack of trust between farmers and traders may also limit the opportunity to switch if some dimension of the contract, such as cocoa quality, is not perfectly observable or enforceable. As emphasized by the large literature on interlinked transactions, trader’s provision of financial services to the farmers may be another important source of differentiation.² In our setting, farmers use advance payments for consumption smoothing or input purchases (e.g. to hire workers that harvest trees).

In the second contribution of the paper, we thus highlight the importance of accounting for credit provision, and more generally interlinked transactions, when studying price transmission and rent distribution in the value chain. We develop two strategies to measure the value of advance payments—one based on baseline correlations in prices and advances across villages and one based on experimental treatment heterogeneity. While using different sources of variation, both methods deliver significant and quantitatively reasonable estimates of the value of advance payments.

of a trader’s inverse supply to her competitors’ quantity and the slope to her own quantity.

²Bardhan (1980), Bell (1988), and Bardhan and Udry (1999) summarize a large body of theory that relates land, labor, output, and credit markets. Blouin and Macchiavello (2017), Casaburi and Willis (2016), Casaburi and Macchiavello (2016), Ghani and Reed (2017), and Macchiavello and Morjaria (2015a) provide primarily empirical contributions.

In the experiment, wholesalers pay treatment traders (which are about 20% of the traders operating in the study region) a per-unit subsidy worth about 5% of the average trader price.³ The subsidy intervention lasts until the end of the harvest season. During the experimental period, treatment traders pay similar prices to farmers than control traders. They are also more likely to provide advance payments to farmers that they report as regular suppliers at baseline (+14 percentage points, a 117% increase). On average, those farmers that did not receive more advance payments from treatment traders (relative to control ones) did receive a higher price from them. We then compute the difference between treatment and control traders in the “effective price,” a price that accounts for advance payments and thus akin to the net present value of the transaction. This difference—which as discussed above is not (generally) a treatment effect—amounts to one-tenth to one-sixth of the subsidy value. By matching this difference to the analog equilibrium equation in the model, we obtain estimates of the differentiation rate that range between 0.1 and 0.2, on a 0 to 1 scale (where 0 implies complete homogeneity and 1 implies that each trader operates as a monopsonist).⁴

In the model, this differentiation rate summarizes the level of competition in the market *for a given number of competitors*. Typically, (quasi-) experimental studies define market boundaries, and thus the number of competitors in a market, ex-ante. Thus, in this standard approach, our experimental results are sufficient to recover the degree of competition. We however acknowledge that defining appropriate randomization clusters may be difficult. For instance, farmers may have the option to sell to traders that operate primarily in other villages, a unit often used as experimental cluster. In addition, traders in many agricultural value chains, including the one we study, buy from farmers in several locations and overlap only partially with the operation areas of other traders. These features complicate again the design of market-level randomizations.

Therefore, in the third contribution of the paper, we propose a methodology to estimate market size, rather than assuming it. For this purpose, we combine the treatment-control difference in prices with quasi-experimental estimates of the pass-through rate from wholesaler to trader prices (i.e. of a common cost shock affecting all traders), which we obtain from plausibly exogenous variation in world cocoa prices. We find a pass-through rate of 0.92, which is in line with evidence from cocoa value chains in neighboring countries (Gayi and Tsowou, 2015). For a given level of differentiation, the model relates the pass-through rate to

³In our setting, traders typically have exclusive relationships and only sell to one wholesaler.

⁴In terms of the model, this implies that the slope of a trader inverse supply to a competitor’s quantity is eighty to ninety percent of the slope to her own quantity.

the “effective market size”. Our estimate of this market size is 14 traders. The interpretation of this result is that traders behave as if the number of their competitors were substantially larger than just the number of traders operating in the same village (7.8).⁵ Overall, our results suggest that the market is quite competitive, with a low degree of differentiation and a large market size. In line with this result, while paying a slightly higher effective price to the farmers, treatment traders purchase substantially more cocoa (+188%) than control farmers during the experimental period. As a benchmark, in a perfectly competitive case, treatment traders would take the entire market as long as they charge an infinitesimally higher price.

The treatment-control difference in prices provides a simple test of the null hypothesis of homogeneous (i.e. non-differentiated) firms. The logic of the test holds in many models of imperfect competition (e.g., oligopoly, monopolistic competition, and search frictions).⁶ However, the estimation of the market structure parameters obviously depends on the specific functional forms of the model, which, as we discuss extensively in the paper, in our case are quite restrictive. We thus provide additional evidence to validate the framework choice. First, since the model is overidentified, we can use different moments (i.e., the *percent* differences between treatment and control traders in prices and quantities) to recover the two market structure parameters. These new parameter estimates are very close to those we had obtained using the first set of moments (we also show that this result is not mechanical). This similarity supports the model choice. Second, we present evidence suggesting that alternative models such as Bertrand, monopsonistic competition, and collusion among traders are not consistent with the data.

Through the model, we can then quantify the impact of the experiment relative to a counterfactual scenario without the experiment. The experimental subsidy raises effective prices by about one-third of the subsidy (i.e., by about 1.5% of the average price level) for treatment traders and by one-sixth for control traders. This confirms that the treatment-control difference in prices, which we find to be one-sixth of the subsidy, is the result of partial price competition between treatment and control traders in response to the intervention. Further, the difference in quantities purchased arises almost entirely from treatment traders stealing from control traders (and from non-study traders, who are about 60% of the traders in the market). Aggregate supply increases at most by 0.9%.⁷ In turn, total farmer

⁵Our baseline empirical approach assumes homogeneous villages in terms of the number of traders per farmer. In the paper, we discuss how to account for heterogeneity along this margin.

⁶The intuition is that in all these models, when firms have heterogeneous costs, there is a unique market price only if the firms are not differentiated.

⁷The lack of change in aggregate quantity is consistent with the fact that the experiment began halfway

revenues increased by 1 to 1.7%. Finally, we quantify the impact of counterfactual subsidy programs treating different shares of traders, given the value of the competition parameters we recovered. We show that, for a given level of pass-through, the welfare impact of a subsidy program targeting only a subset of traders varies with the differentiation rate. Given that industrial policies often target only a subsample of firms, this result emphasizes the importance of estimating specifically the differentiation rate parameter.

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. It relates to a growing body of work that seeks to estimate the equilibrium effects of (quasi-)experiments using market-level randomization.⁸ Recent examples include Crépon et al. (2013), Cunha et al. (2015), Lalive et al. (2015), Baird et al. (2014), Hildebrandt et al. (2015), Burke (2014), Mobarak and Rosenzweig (2014), Muralidharan et al. (2016), McKenzie and Puerto (2017), and Breza and Kinnan (2016).⁹ Several recent papers that study competition using market-level variation are particularly related to this paper (Busso and Galiani, 2014, Jensen and Miller, 2016, Mitra et al., 2013, Rotemberg, 2017, Bergquist, 2017). In particular, Rotemberg (2017) identifies the spillover effects on larger firms of a subsidy available only to small firms in India, using variation in exposure to eligible competitors, and Bergquist (2017) randomizes subsidies across markets in Western Kenya to maize consumers and traders, in order to infer market structure parameters of the Atkin and Donaldson (2015) framework.

Relative to this body of work, our paper provides a proof of concept of how purely individual-level randomization can shed light on market structure parameters even in the presence of competitive spillovers. It complements the contribution of Atkin and Donaldson (2015) by using exogenous cost shocks that vary across firms in the same market, rather than across markets, to inform the measurement of competition. As a first step in this direction, the paper leaves several avenues open for future work. For instance, since the experiment introduces heterogeneity across subsidized and unsubsidized firms, the Atkin and Donaldson (2015) general model of symmetric imperfect competition is not readily applicable.¹⁰ We

in the harvest season and thus farmers had limited options to increase their output volume (for instance by harvesting marginal fruits, at the top of the trees, or by reducing processing losses).

⁸Other 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. (2014), and Emran et al. (2017).

⁹Researchers have started acknowledging the need to expand the size of the cluster beyond the village. For instance, Muralidharan and Niehaus, 2016 discuss the benefits and challenges of randomization “at scale”.

¹⁰Weyl and Fabinger (2013) illustrates the complications that arise when modeling asymmetric firms in the framework.

thus work with a less general model (which we however attempt to validate). With this in mind, we also highlight several benefits. First, as discussed, individual-level randomization may be the only option if there are economic constraints or structural limitations (e.g., if the program targets only one or few large markets). Second, the differentiation rate among firms, the parameter we recover solely from the experiment results, is policy relevant. It shapes, for a given pass-through, the distributional impact of subsidies that target a subset of the firms in the market. Third, in principle, individual-level randomization can generate distinct market structure estimates in each location. Future research may explore synergies between market-level and individual-level experimentation to understand market structure.

Finally, the paper contributes to the literature on the competitiveness of agricultural markets in developing countries. Our experimental design provides a simple test that we hope can be replicated in other settings. This approach complements previous studies of competition that have primarily relied on observational data, analyzing trader price-cost margins (for the case of Sub-Saharan Africa, see, e.g., Fafchamps et al., 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). Our evidence of competitive crop markets is in line with the review by Dillon and Dambro (2016). We also add to this evidence by showing that low differentiation and high competition can coexist with high prevalence of financial service provision (e.g. credit) from traders to farmers. This insight may be useful in considering the external relevance of our findings to other settings.

The rest of the paper proceeds as follows. Section 2 provides background information on the study setting and the experimental design. Section 3 presents the model and the strategy to recover the competition parameters. Section 4 presents the experimental results and recovers the trader differentiation rate. Section 5 combines our experimental results and quasi-experimental estimates of pass-through to further recover market size. Section 6 provides additional evidence in support of the model choice and it presents model-based analysis of the impact of the experiment. Section 7 concludes.

2 Study Setting and Experimental Design

2.1 The Sierra Leone Cocoa Value Chain

Though Sierra Leone accounts for only a small share of the world production, cocoa is important nationally. The crop comprised 8.6% of exports in 2009, and is the largest export crop by value, according to UN COMTRADE. The harvest season typically lasts from August until the end of the year. Climatic differences cause variation in specific harvest times both across locations and across years in a given location. A given farmer may be harvesting at different times depending on the location of her plots or the age of the trees.

The within-country cocoa trade in Sierra Leone is fragmented across many traders, and the supply chain has many links, similar to other agricultural markets in developing economies (for examples in Africa see Fafchamps et al., 2005 and Osborne, 2005).¹¹ Farmers sell to traders, who sell to wholesalers in small towns, who in turn sell to exporters in larger towns, who in turn sell to buyers at the port. While it is important to study the degree of competition in each of the links of the supply chain, we focus on the link closest to production, and leave the examination of other levels for future research.

As emphasized by Atkin and Donaldson (2015), when looking at prices, it is important to focus on narrowly defined homogeneous goods. The quality of cocoa is indeed 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 measure of quality in the market, wholesalers and traders agree on broad determinants of quality that are consistent with international standards (see Fold, 2005). A quality premium exists in the market to some extent. In our experiment, in order to ensure we measure prices for a well-defined homogeneous good, we worked with the partner wholesalers to develop a quantitative quality grade that correlates well with baseline prices. When traders arrive at the warehouse, inspectors hired by the research team sampled 50 beans from each bag, and used them to create an index of quality—grades A, B or C—which was then applied to each bag.¹² The analysis in this paper focuses on grade A cocoa, the grade targeted by the experimental subsidy, unless otherwise specified.

¹¹Sierra Leone’s cocoa industry is similar to those in Cameroon, Ivory Coast, and Nigeria, all of which, unlike Ghana, liberalized during the 1990s and became similarly fragmented (see, e.g., Gilbert et al., 2009).

¹²Appendix B provides details.

2.2 Experimental Design

We developed our experiment in partnership with five privately owned wholesalers in Sierra Leone’s cocoa producing Eastern Province, in the towns of Segbwema, Pendembu, and Kailahun. These wholesalers collect cocoa in their warehouses use a network of trader with whom they have exclusive relations (i.e. a trader almost always delivers cocoa only to one wholesaler). Traders purchase cocoa from farmers and within a few days deliver to wholesalers. These then sell to exporters in the provincial capital of Kenema or in Freetown. Our sample includes 80 traders, henceforth *study traders*. This comprises almost all of the traders who do business regularly with these wholesalers.

During the experiment, a random subset of 40 traders received a bonus of 150 Leones —5.6% of the average wholesale price —when selling grade A cocoa purchases from farmers to the wholesalers. The experiment ran from mid-October to the end of December of 2011, roughly the end of the harvest season. At the beginning of the experiment, traders were informed the treatment would last until about the end of the harvest season.

Randomization occurred at the trader level. We implement a pairwise randomization strategy (see, e.g., Bruhn and McKenzie, 2009): we matched traders within wholesalers according to a self-reported estimate of the volume of purchases since the beginning of the cocoa season and then assigned treatment and control within pairs.¹³

2.3 Data Collection and Summary Statistics

Over the course of the experiment, we collected a variety of data from traders. At baseline, we interviewed each trader about his experience in the industry, and collected basic demographic indicators. We also asked traders to list each farmer they buy from regularly and all of the villages in which they buy. Table 1, Panel A presents summary statistics and shows that 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.¹⁴ 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

¹³Of 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.

¹⁴Figure 1 presents a map of the study setting.

(i.e. working with other wholesalers). In Section 3.4, we discuss the implications of this fact for our estimation approach.

The provision of loans by traders to farmers is an important characteristic of this industry.¹⁵ Traders offer to purchase cocoa in advance before and during the harvesting season. Farmers use the advance payments for production (e.g., hiring workers for harvesting) or for consumption smoothing. Farmers then pay for these advances by selling at a below market price for subsequent sales.¹⁶ Contracts are enforced through customary authorities (see, e.g., Acemoglu et al., 2014 and Sandefur and Siddiqi, 2013) or through relational contracts (see, e.g., Fafchamps, 2003 and Macchiavello and Morjaria, 2015b). Study traders report having given at least one loan to about 70% of the suppliers listed at baseline in the previous year.

During the experiment, when traders arrived at the warehouse, inspectors from the research team measured and documented the quality of their shipment.¹⁷ Enumerators then asked traders the price per pound they paid to farmers and the name of the village where the cocoa mostly originated. Traders typically 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.¹⁸ In addition, to study the impact of trader treatment on advance payment provision, in November and December we asked again the traders if they had given loans in the previous month to the farmers listed at baseline.

In the three weeks preceding the intervention, 56 of the 80 traders (27 control and 29 treatment) visited the warehouses. Table 1 Panel B shows that treatment and control groups are balanced along volume purchased and prices paid to farmers. During the experiment, 74 traders visited the warehouse (36 controls and 38 treatment).¹⁹

¹⁵Other agricultural value chains feature supply chain credit provision. Emran et al. (2017) study the impact of such credit provision on the price response to policy reforms.

¹⁶Interviews to farmers and traders suggest that these contracts do not define the price of the transaction, which is instead based on the market price at the time of the delivery. Rather, they define the amount to be deducted from the final payment (i.e., the interest).

¹⁷Data collection ran from September 24th, 2011 to December 31st, 2011. The intervention started on October 15th, 2011. Because of project budget constraints, data collection was suspended for approximately two weeks and half between late November and early December.

¹⁸In Section 4, we discuss potential concerns arising from the self-reported nature of the data.

¹⁹The regressions presented in the rest of the paper include pair fixed effects. Therefore, we effectively drop pairs of those traders that did not visit the warehouses.

3 A Simple Model of Buyers' Imperfect Competition

This section presents a simple framework of oligopsonistic competition among buyers and it models the impact of a subsidy to a subset of buyers (akin to our experimental treatment). We derive closed form solutions for equilibrium prices and quantities of treatment and control traders and for the differences in outcomes between the two groups.

3.1 Setup

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 inverse supply buyer i faces in a village is:²⁰

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

We adapt the standard model of linear demand and differentiated producers (see, e.g., Vives, 2001) to our setting, which features imperfect competition among buyers.²¹ Section 3.3 discusses the assumptions of the framework (including the linearity of supply curve) and Section 6.1 presents evidence to validate the model choice.

Following the literature, we define the *differentiation 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.²² In the model, differentiation allows in a reduced form for potential unobserved aspects of a specific buyer-seller relationship that may give the

²⁰This inverse supply can be microfounded by assuming a representative producer whose cost function features *love for variety*. Specifically, the producer profit function is: $V(p_1^E, \dots, p_n^E, q_1, \dots, q_n) = q_0 + \sum_{i=1}^n p_i q_i - C(q_1, \dots, q_n) = q_0 + \sum_{i=1}^n p_i q_i - (\alpha \sum_{i=1}^n q_i + \frac{1}{2} \beta \sum_{i=1}^n q_i^2 + \gamma \sum_{j \neq i} q_i q_j)$, where q_0 is the output that is not sold to traders (e.g., consumed, not harvested) and q_i is the output sold to trader 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.

²¹As we discuss in Section 4.2, p_i , the actual value paid to the farmer, may combine payments made at different times (e.g., pre-harvest advances and post-harvest payments), which have different values for the producers.

²²The direct supply function is $q_i = a + b p_i - c \sum_{j \neq i} p_j$, with $a \equiv \frac{\alpha}{\beta + \gamma(n-1)}$, $b \equiv \frac{\beta + \gamma(n-2)}{(\beta + \gamma(n-1))(\beta - \gamma)}$, $c \equiv \frac{\gamma}{(\beta + \gamma(n-1))(\beta - \gamma)}$.

buyer market power. As discussed in Section 1, differentiation can originate from a range of reasons, including search costs, trust, repeated relationships, and advance payment provision. We show that accounting for interlinked transactions (e.g. advance payments) is important to understand differentiation among traders. However, we do not aim to provide a complete breakdown of the relevance of these individual sources of differentiation.

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 (net) resale price, and p_i is the price paid to producers.²³

We assume buyers are ex-ante symmetric in the resale price v .²⁴ 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).

We assume *Cournot oligopsonistic competition*: each buyer sets quantities strategically, taking into account competitors' choices. In Section 6.1, we provide evidence in support of this choice and discuss alternatives such as Bertrand competition, monopsonistic competition, and trader collusion.

3.2 Equilibrium

We consider a “group-symmetric” equilibrium in which firms in the same treatment group behave similarly.

²³A given buyer can compete in multiple villages. However, the choices she makes across villages are assumed to be independent. Thus, we restrict the analysis to the village level and omit the village index in the equations above.

²⁴Appendix A.1 relaxes this assumption.

Equilibrium

Using standard algebra, we can derive the quantities set by treatment and control buyers:

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

From the inverse supply functions in Equation 1, we then obtain equilibrium prices:

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

It can be shown that 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 suggest that treatment changes the behavior of control firms, as well as treatment ones. In experimental terms, the strategic response of unsubsidized traders to the subsidy of their competitors represents a violation of the Stable Unit Treatment Value Assumption (SUTVA).

Recovering Trader Differentiation from the Experimental Results

The differences in equilibrium outcomes between treatment and control traders are:

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

and

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

Observe that 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: control traders can pay a price lower than do treated traders.²⁵ On the other hand, Δq is *decreasing* in Γ (for given β): if traders are homogeneous, the treatment traders can expand

²⁵If $\Gamma = 1$, i.e., each buyer is a local monopsonist, the linear supply function implies that each monopsonist passes through one-half of the subsidy, that is $\Delta p = s/2$.

by taking market share from control traders. Both the price and the quantity differences are increasing in the value of the subsidy, s , but do not depend on the share of treatment traders, μ . This is because an additional treatment trader takes away quantity from both control traders and other treatment traders, thus the impact on the difference between the two types of traders is ambiguous. In the case of the specific functional form we adopt, this impact is zero. By matching the theoretical expression in Equation (5) to its empirical analog (which we derive in Section 4.1.2), we can recover the differentiation parameter, Γ .²⁶ If one is willing to assume market size, as standard experimental and quasi-experimental approaches do, the experimental estimate of the treatment-control difference in prices is sufficient to estimate the degree of competition among traders, since this depends only on Γ . This approach may be useful when the researcher is confident about plausible market boundaries, but cluster-level randomization is unfeasible for power, logistics, or budgetary reasons.

Recovering Market Size: Combining Experimental and Quasi-Experimental Results

The model also delivers a solution for how buyer prices respond to a market-level shock in the resale price, v , common for all traders. The (constant) pass-through rate is

$$\rho \equiv \frac{\partial p_C}{\partial v} = \frac{\partial p_T}{\partial v} = 1 - \frac{1}{1 + \Gamma + n(1 - \Gamma)}, \quad (7)$$

which is decreasing in Γ and increasing in n . Equation 7 shows that, for given Γ , the pass-through rate allows us to recover the market size parameter, n .

3.3 Discussion of the Assumptions

The general model of symmetric imperfect competition of Atkin and Donaldson (2015) is not readily useful to interpret heterogeneous firms pricing in response to asymmetric taxes/subsidies within the same market (i.e. our key experimental moment).²⁷ To relate the experimental results to the theory, we then work with a simpler model, which makes

²⁶We note that this estimation procedure does not use the treatment-control differences in quantities (Equation 6). This is because Δq depends on the level of β , as well as on Γ . In addition, using this moment would require a definition of market size, which we later take as a parameter to estimate. Section 6.1 presents a different method that uses *percent* treatment-control differences in prices and quantities. The use of percent differences does not require an *ex ante* definition of market size.

²⁷Refer also to the discussion of the general model with asymmetric firms in Weyl and Fabinger, 2013.

several additional assumptions. Here, we discuss some of the key ones. First, the model is static. This may be missing important features of the economic environment we work in. For instance, advance payment provision, which in our setting turns out to be an important response margin for treated traders, hinges on repeated interactions.²⁸ Second, we rely on specific functional forms. For instance, we focus on linear supply, rather than working with unrestricted supply elasticity.²⁹ Similarly, we assume that v and p do not depend on quantities, thus ruling out non-linear pricing and other non-constant trader marginal costs.³⁰ Third, agents are symmetric, aside from the heterogeneity introduced by the experiment (i.e., the experimental subsidy, s). Fourth, we assume the experiment does not change the competition structure. This could happen, for instance, if the subsidy leads to entry or exit.³¹

Because of these restrictive assumptions, it is important to validate the choice of the model. We do this in Section 6.1. First, in the spirit of an overidentification test, we show that different moments of the model lead to similar estimates of the market structure parameters. Second, we show that data do not seem to support alternative models. Third, we discuss how our estimation framework can accommodate certain forms of baseline heterogeneity among traders. We leave to future research to generalize the framework by relaxing the other assumptions discussed above.

3.4 The “Theory Experiment” and the “Field Experiment”

The field experiment setting obviously presents some deviations from the stylized environment of the model. Here, we discuss how these discrepancies may affect our estimates of the trader differentiation rate. First, we started the experiment in the middle of the harvest season. It is possible that, by that time, traders had already locked in purchases from some farmers with advance payments. Thus, the degree of differentiation may have been lower if we had started the experiment before the harvest season. Had the wholesalers announced the subsidy earlier, it is possible that treatment traders may have accessed an even larger pool

²⁸We discuss this topic further in Section 4.2.

²⁹Among other reasons, this 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 and Armstrong and Vickers, 2015).

³⁰For instance, Attanasio and Pastorino (2015) presents evidence of nonlinear pricing in rural Mexican villages and proposes a model of price discrimination to account for this nonlinearity.

³¹In another example, the subsidy could relax liquidity constraints of treated traders, thus relaxing their potential quantity constraints and changing the extent to which they can compete. Furthermore, subsidies could foster trader investment in transport costs or other technologies (for a study of the relation among industrial policy, competition, and innovation, see Aghion et al., 2015).

of contestable farmers.³² In turn, our estimate of the differentiation rate, which as we will see is already quite low, may be an upper bound relative to that obtained in a season-long experiment.

Second, 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 should 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 5.1 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.

Third, another important distinction between the basic model presented in Section 3 and the experimental setting concerns the presence of non-study traders. These comprise about 60% of the traders operating in the study region. In principle, these could be different from the study traders (control and treatment) at baseline. Importantly, the model presented in Section 3 is robust to the presence of such traders. Specifically, in Appendix A.2, we discuss a 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 . We show that the equilibrium treatment-control differences, Δp and Δq (Equations 5 and 6), and the pass-through rate ρ (Equation 7) are unchanged.³³ Therefore, the estimation approach is robust to the inclusion of non-study traders.

³²As we discuss in Section 4.2, our static framework does not model the trader choice of raising prices vs. increasing advance payment and thus we can only speculate on how the timing of the intervention would affect the relative intensity of the trader response along these two margins.

³³In this augmented model, n is the total number of traders, i.e., study and non-study.

4 Experimental Results and the Estimation of the Trader Differentiation Rate

In this section, we first report treatment-control differences in trader prices, advance payment provision, and quantities purchased during the intervention period. Second, we quantify the value of advance payments and thus the treatment-control difference in *effective prices*, akin to the net present value of the payment. Third, we present estimates of the differentiation rate.

4.1 Experimental Results

We report differences in outcomes between treatment and control traders. As discussed above, these differences cannot be interpreted as treatment effects in the standard potential outcomes framework of Rubin (1974): as the model in Section 3 clarified, the subsidy affects the behavior of both treatment and control traders, as they compete for the same suppliers. However, in our approach, these differences are nevertheless crucial estimation objects. By matching them to the theoretical counterparts derived in the previous section, we can later recover the competition parameters of the model.

4.1.1 Prices

First, we focus on prices that traders pay to farmers. Enumerators asked traders the purchase price for each shipment. If the traders made payments at different times (e.g., before and after harvesting), enumerators recorded the total value traders paid for cocoa, not just the harvest one. We denote this variable with \tilde{p} , so to differentiate it from the *effective price*, p , we focused on in the model. As discussed above, effective prices take into account the relative values of payments made at different times.

Figure 2 displays the price results graphically. It shows weekly averages for:³⁴ i) world prices (right y axis);³⁵ ii) wholesaler prices; iii) trader prices paid by treatment traders;

³⁴In principle, dispersion in prices for transactions on identical cocoa would also be informative of trader differentiation. In practice, cross-sectional variance in prices is also likely to reflect measurement error and variation in transaction characteristics, such as date, exact location, transport costs, or specific product features. In our baseline price data, the coefficient of variation is 0.07.

³⁵Specifically, 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 an exchange rate of 1 USD=4,400 Leones.

iv) trader prices paid by control traders. The vertical red line marks the inception of the intervention period. The graph displays two key features. First, trader prices follow closely wholesaler prices and these move with world prices. In particular, domestic prices respond to the sharp decrease in the world price that occurred in November 2011. Second, this preliminary graphical analysis displays no obvious gap in prices that treatment and control traders pay to the farmers.

We estimate the following regression, where an observation is a shipment s delivered by trader i of randomization pair z in week t :

$$\tilde{p}_{sitz} = \eta_z + \eta_t + \pi^{\tilde{p}} \text{Treat}_i + \epsilon_{sit}, \quad (8)$$

where η_z and η_t are randomization pair and week fixed effects, respectively. We cluster standard errors at the trader level (the unit of treatment).³⁶

The term $\pi^{\tilde{p}}$ is the coefficient of interest. Table 2, Column (1) presents a regression without week fixed effects: the coefficient is $\hat{\pi}^{\tilde{p}} = -32.5$ (s.e.=47.2). In Column (2), we introduce week fixed effect and the coefficient becomes -5.5 (s.e. = 14.9). While the two coefficients are not statistically distinguishable from each other or from zero, that the coefficient is higher in absolute value without week effects suggests that selection in when to sell matters. In particular, it appears that the experiment induced treatment traders to stay longer in the market at the end of the season, when prices were lower.

One additional concern is that the treatment may induce selection in which traders make purchases and in which locations traders visit. For this reason, in Columns (3)-(5), we add controls referring to the trader and to the village where the majority of cocoa in the shipment originated.³⁷ The coefficient of interest is quite stable when including these controls. This suggests that the selection concerns described above cannot drive the results.³⁸

Overall, the various specifications provide evidence that prices did not differ between treatment and control traders. As highlighted by the model, this does not necessarily imply that traders did not respond to the subsidy. Rather, it may reflect the response of both treatment and control traders to the subsidy. Suggestive evidence in line with this hypothesis

³⁶Results are similar when allowing for double clustering by trader and village (Cameron et al., 2012).

³⁷See notes to Table 2 for a list of controls. Eighty 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.

³⁸In results not presented, we also tested for effects on the prices of B and C grade cocoa. We find a statistically significant difference for grade B prices (the point estimate is 37, which is still very far from the value of the subsidy). Field interviews suggest that treatment traders were somewhat more willing to pay the grade A price for cocoa that had some probability of being grade A.

comes from the fact that the treatment-control price gap is larger in the first weeks of the experiment (e.g., in the first three weeks of the experiment, the treatment coefficient is 31, $p=0.09$), and then decreases. In Section 6.2, we use the model and our estimates to quantify the impact of the experiment on prices paid by treatment and control traders (relative to a counterfactual without the experiment).

4.1.2 Advance Payments

To investigate the treatment-control differences in the provision of advance payments during the intervention period, we estimate the following regression:

$$AdvancePayment_{fiz} = \eta_z + \pi_1^a Treat_i + \nu_{fi} \quad (9)$$

An observation is a farmer listed as a regular supplier in the trader baseline. $AdvancePayment_{fiz}$ is an indicator of whether farmer f received an advance from trader i of randomization pair z during the course of the experiment.

Table 3 presents estimates of π^a , the coefficient of interest. In Column (1), we run a linear probability model where the outcome is a dummy equal to one if credit was provided to a farmer. The difference between treatment and control traders is substantial: farmers reported as regular suppliers by treatment traders in the baseline listing are 14 percentage points more likely to receive credit from these traders, relative to a control mean of 11 percent.³⁹ Columns (2)-(4) show that results are similar when adding trader controls, village controls, and both set of controls. These results suggest that while there were not differences in trader prices, there was likely difference in effective prices paid, once the value of payments being provided in advance is taken into account. We will address this point in Section 4.2 below.

One important limitation of our analysis is that, due to funding constraints, price and credit data are self-reported by traders at the time of their visit to the warehouses. One may be concerned subsidized traders may over report the price paid and credit offered to farmers. Two observations mitigate this concern. First, we note that, if true, this would bias *upward* our estimate of the differentiation rate, which, as discussed, we find to be quite low. Second, simple reporting bias does not easily explain a positive treatment-control difference for advance payments but not for prices.

³⁹This control mean refers only to short-term credit given during the 2.5 months of the experimental period. It is not inconsistent with the observation that, in the twelve months before the experiment, traders reported giving at least one loan to approximately 70% of their suppliers, as reported in Section 2.3

4.1.3 Quantities

Finally, we investigate the treatment-control difference in quantities purchased during the experiment. Figure 3 shows the weekly amount purchased by the study traders and then by treatment and control groups separately. Several patterns emerge. First, purchases of treatment and control are balanced in the three weeks before the intervention. Second, throughout the intervention, treatment traders purchase substantially higher volumes than control ones. Third, total quantity purchased by study-traders increases after the beginning of the experiment. This observation is consistent with the idea that treatment traders may have gained market shares at the expense of non-study traders, as well as of control traders. Finally, toward the end of the experiment, there is a stark reduction in total quantities purchased, consistent with field reports that the season was essentially over by that time.

Table 4 presents regression results from the following regression models:

$$Quantity_{izt} = \eta_t + \eta_z + \pi^q \text{Treat}_i + \zeta_{it} \quad (10)$$

where an observation captures the total purchases of cocoa trader i of randomization pair z in week t . During the experiment, treatment traders on average purchase 527 pounds per week more than control traders, roughly a 188% difference.⁴⁰ The results are robust when including trader controls in Column (2). Overall, this is a large impact of the treatment.⁴¹ Given that farmers had limited opportunities to increase production by the time the intervention started, it seems likely that market-stealing effects may drive the results. Treatment traders could steal from both control traders (20% of the market) and non-study traders (60% of the market). The analysis in Section 6.2 supports this conjecture.

⁴⁰Consistent with the large difference in quantities purchases, 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.

⁴¹We include all the eighty traders in the sample, assigning value zero to trader-week pairs with zero purchases, including for the traders that never showed up during the experimental period (results are similar when dropping these traders). The treatment-control difference on quantities is substantially smaller (in absolute value) in the last three weeks of the experiment. On the other hand, there is no significant difference across these periods in the price regression described in Section 4.1.1.

4.2 The Value of Advance Payments

We now quantify the value buyers provide to farmers through advance payments and then estimate the treatment-control difference in the *effective price*. As discussed above, the effective price is akin to net present value and thus reflects the relative value of payments made at different times. This treatment-control difference in the effective price is what will be matched to the model’s moment in order to estimate trader differentiation.

It is important to emphasize that our framework does not model the trader’s choice to pass value to the farmer through advance payments or through higher prices.⁴² Moreover, as discussed in Section 3.3, the static framework obviously does not capture the repeated game nature of advance contracts. Accounting for these elements would require modeling a repeated game framework of trader competition featuring multiple choice variables for the traders. Such an approach may not generate easily closed-form solutions for the treatment-control differences and it may feature multiple equilibria, thus substantially complicating the estimation. We therefore assume that traders face a separable problem. First, they set their effective price based on the inverse supply 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 only the effective price (the net present value of the payments), not its breakdown.

While we do not explicitly model the trader choice to provide advances, we nevertheless need to measure the value of these advance payments, and thus the effective price. For this purpose, we define the effective price paid by buyer i as it follows:

$$p_i = \tilde{p}_i + \lambda \cdot ShareAdvances_i, \quad (11)$$

where \tilde{p}_i is the total monetary amount paid by the trader, $ShareAdvances_i$ is the share of farmers to whom trader i provides advance payments.⁴³ Therefore, λ is the extent to which farmers value advance payments (i.e., the rate of substitution of the indifference curve between prices and advances).⁴⁴ We propose two simple strategies to estimate the value of

⁴²We can only speculate this could be due to the fact that treatment traders can use some of the extra profit to secure future supply with advance payments and that control traders do not have sufficient funds to compete along this margin.

⁴³Ideally, our data would include the specific amounts of payments made at different times. However, we only observe shipment-level prices and an indicator of whether a certain farmer receives advance payments.

⁴⁴Since the effective price enters both the farmer utility function and the trader profit function, we are assuming that farmers and traders have the same rate of substitution, λ .

advance payments.

Approach 1: Cross-Sectional Baseline Correlations

First, we infer the value of advance payments from the relation in the baseline cross section between shipment prices (i.e. the total monetary amount paid by the trader for a given shipment) 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.⁴⁵ Table 5, column (1) shows that moving from a village where no farmer receives advance payments at baseline to a village where each farmer receives advance payments decreases the amount of total payments paid by the trader by 149.6 Leones (s.e. = 74.6).⁴⁶ An interpretation of this result is that a farmer is indifferent between a trader paying a certain price and another trader who pays 149 Leones less but provides advance payments. The result is robust to the inclusion of village-level controls (column 2). These results, while based on limited baseline data, provide initial evidence consistent with the fact that farmers value advance payments and are willing to accept lower prices from traders for this service.

Approach 2: Heterogeneity in Treatment-Control Differences

A second approach is to infer the value of advance payments from the relationship between the treatment-control differences in prices and those in advance payments. In the model, traders respond to the treatment by increasing the *effective price* by a certain amount. This response can come in the form of higher prices or more frequent advance payments, and the farmer’s indifference curve between the two will define which pairs of adjustments provide equivalent value. The slope between the two response margins identifies their relative value, or how much less a trader who increases her advance payments needs to adjust her prices.

For this purpose, we modify Equation (9) to allow for heterogeneity in the treatment-

⁴⁵“Village-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 (see Section 4.3). 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).

⁴⁶For this regression we only use data from the subset of villages in which we have pre-treatment data.

control differences across villages and trader characteristics:

$$AdvancePayment_{fizv} = \eta_z + \pi^a(\text{Treat}_i) + (\text{Treat}_i \times X'_v)\pi_v^a + X'_v\beta_v + (\text{Treat}_i \times X'_i)\pi_i^a + X'_i\beta_i + \nu_{sipv}, \quad (12)$$

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:

$$\tilde{p}_{sizvt} = \eta_z + \eta_t + \pi^{\tilde{p}}(\text{Treat}_i) + \pi_a^{\tilde{p}}(\widehat{DTC}_{iv}^a \cdot \text{Treat}_i) + X'_i\beta_i + X'_v\beta_v + \epsilon_{kiptv}. \quad (13)$$

If total payments and advance payments are substitutes (i.e., $\tilde{\lambda} > 0$), then $\pi_a^{\tilde{p}} < 0$.⁴⁷

Figure 4 provides some intuition for how this procedure works, and 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 institutions.⁴⁸ The scatter displays a negative relation: the regression line has a slope of -271. Table 6 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 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), \widehat{DTC}_{iv}^a is predicted using only chiefdom dummies. The estimate using these dummies predicts that a village where treatment traders are 14 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 47.8 Leones lower than a village with no difference in advance payments. This is economically relevant as it accounts for a reduction in the treatment difference of about one-third of the subsidy

⁴⁷Since \widehat{DTC}_{iv}^a is collinear with the vector of controls, its level is not included in the estimating equation.

⁴⁸Unfortunately, our data do not include explicit information on contract enforcement institutions and thus we cannot study heterogeneity in the trader responses by this variable.

value. 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 trader covariates. While the magnitude of the coefficients falls across columns, the core result holds: price and advance payment responses are substitutes.

Discussion

One may worry that villages with higher shares of advance provision (in either baseline level or treatment response) may have some other feature that lowers prices and thus that we are capturing a spurious relationship. While we cannot completely rule out this concern, three observations mitigate it. First, both approaches derive significant, qualitatively similar, and quantitatively meaningful results even if they use different variation. In particular, the baseline provision of advances (i.e., the source of variation used for approach 1) is not a major factor in the heterogeneity used in approach 2.⁴⁹ Second, we control for a range of village and traders covariates.⁵⁰ Third, assuming that the loan covers the entire purchase and that the loan duration is two months (one month), the implied interest rate is around 3% (6%) per month. This is a high rate, but not inconsistent with prevailing interest rates.⁵¹

Treatment-Control Differences in *Effective Prices*

Having obtained estimates of the value of advance payments — λ in the model—, it is then possible to estimate the treatment-control differences in the effective price, p : $\hat{\pi}^p \equiv \hat{\pi}^{\bar{p}} + \hat{\lambda} \cdot \hat{\pi}^a$. Based on Tables 2 and 3, we specify $\hat{\pi}^{\bar{p}} = -6.9$ and $\hat{\pi}^a = .14$. The values of λ are 150 when using the baseline correlations and 210 when using the treatment heterogeneity (we use the specification with all the interactions, as presented in Table 6, col. 3). The two approaches then lead to point estimates for the treatment-control difference in effective prices of 14.1

⁴⁹We note that the estimates of λ derived in the second approach (210-334) are generally larger than the estimates from the cross-sectional analysis in Table 5 (147-150). However, the baseline advance payment variable, which we use in the first approach, captures whether the farmer had received advances in the twelve months before the baseline, a longer time horizon than the one of the experiment (two months). Thus, the two advance payment dummies may capture different intensities of advance payments.

⁵⁰In principle, it is also possible that some of this credit is non-interlinked. In response to higher margins, traders could invest part of the extra profits in their lending business, without a link to the trading activities. However, qualitative evidence from interaction with the traders suggest that traders use the loans to secure supply and that the bulk of the credit is repaid through lower prices at harvest time.

⁵¹According to the World Development Indicators, the average lending interest rate in the last fifteen years was between 21% and 25% per year. In the inventory credit evaluation described in Casaburi et al. (2014), rates on subsidized collateralized loans for agricultural smallholders were 22% per year. Rates on uncollateralized agricultural loans and on moneylender loans are likely to be to substantially higher.

and 22.7, respectively. These results suggest that, during the intervention, treatment traders pay an effective price higher than the control price by 10% to 15% of the subsidy value (150 Leones). This result also suggests that, on average, those farmers that did not receive a differential increase in advance payment from treatment traders (relative to control ones) did receive a differentially higher price from them. For these farmers, according to our estimates, the average difference in prices is one sixth of the subsidy (i.e. equal to the difference in effective prices). Finally, we emphasize that, as discussed above, the difference in effective prices could arise from price competition between traders. Therefore, it does not (generally) measure the impact of the experiment on treatment traders (relative to a counterfactual without the experiment).

For the approach based on baseline correlations, we compute 90% confidence interval by jointly estimating $\hat{\pi}^p, \hat{\lambda}, \hat{\pi}^a$. We obtain [-10.93, 37.53] and we can reject the null that the difference in effective price is zero at $p=0.23$. For the second approach based on treatment heterogeneity, we instead use bootstrap since $\pi_a^{\tilde{p}}$ is a generated regressor. We cluster resampling at the randomization pair level. The confidence interval is [-31.11, 61.33].

4.3 Recovering the Trader Differentiation Rate

To recover Γ , we match Equation (5) to its empirical analog:

$$\Delta p \equiv \frac{s\Gamma}{1 + \Gamma} = \hat{\pi}_1^p + \hat{\lambda} \cdot \hat{\pi}_1^a \quad (14)$$

Before presenting the results, we discuss an additional assumption that is necessary to reconcile the theory to the available data: we assume that advance payments to *regular* suppliers—for whom we have data on advances—are representative of advance payments to all the farmers. Unfortunately, it is not clear in which direction a violation of this assumption would 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.

We solve Equation (14) for Γ using the two values of the treatment-control differences in effective prices that we derived from the two methods in Section 4.2: $\lambda = 149$ and $\lambda = 210$. We obtain the following results for point estimates and 90% confidence intervals: *i*) with $\lambda = 149$, Γ is 0.10 [-.10, .29]; *ii*) with $\lambda = 210$, Γ is 0.176 [-.18, .63].⁵²

⁵²When we ignore that λ is a generated regressor, we can reject $\Gamma = 0$ at $p=0.13$. If we ignored advance

5 Combining Experimental and Quasi-Experimental Results to Estimate Market Size

Our experimental results are sufficient to describe competition if, as it is standard in experimental approaches, one assumes market size (i.e. the size of the cluster). In this section, however, we show that, by combining experimental and quasi-experimental estimates of the pass-through rate, one can estimate market size, rather than assuming it.

5.1 Pass-Through of Industry-Wide Price Shocks

We first study how trader prices respond to common changes in wholesaler prices (w in the model). As preliminary evidence of high responsiveness, Figure 2, discussed in Section 4.1.1, showed a stark reduction in prices (around 22%) in the final month of the experiment, following a reduction in world prices and wholesaler prices. Table 7 presents the results of the regression analysis. Column (1) reports a basic OLS regression of the trader price vs. the wholesaler price with trader fixed effects. Throughout the table, standard errors are clustered by date.⁵³ The coefficient estimate is 0.89. This initial result suggests a high-level of pass-through from traders to farmers (and low trader markdowns).

The change in wholesaler prices may be correlated with local supply shocks. To address this concern, we instrument wholesaler prices with the international price of cocoa, as measured 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. The instrument leads a very strong first stage (Kleibergen-Paap F-stat=14,024). Column (2) in Table 7 shows that the pass-through rate estimate is 0.92. In Column (3), we also add month fixed effects. The coefficient is stable. Finally, Column (4) shows that the coefficient is robust to the inclusion of village fixed effects (though the first stage becomes substantially weaker).⁵⁴ Overall, our results suggest high level of pass-through. This is consistent with Gayi and Tsowou (2015), who shows 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 of around 90%.⁵⁵

payments, the point estimate of Γ would be -0.03.

⁵³The pass-through estimates remain significant at $p < 0.01$ if we cluster by week.

⁵⁴We also obtain similar results when using a lag of international prices (of a day or a week) in our regression.

⁵⁵High pass-through in these countries also suggests that our results are not driven by the self-reported

5.2 Recovering Market Size

By matching Equation (7) to its empirical analog, and by using our previous estimate of Γ , we can now recover n , the number of traders competing in a given location:⁵⁶

$$\rho \equiv 1 - \frac{1}{1 + \Gamma + n(1 - \Gamma)} = \hat{\rho} \quad (15)$$

Before presenting the results, we discuss a caveat in the interpretation of the results. As discussed in Section 4.2, we did not collect high frequency (i.e. transaction-level) data on advance payments. Therefore, we cannot study how, on a day-to-day basis, advances respond to changes in the industry-wide price. We assume advance payments do not respond to the transitory world price shocks and thus that the price pass-through we measured in Section 5.1 is equal to the pass-through of the effective price. While this seems a plausible assumption given the transitory nature of the world price shocks, we cannot test it. However, we note that omitting any potential credit component of pass-through would mean that our estimate of n is a lower bound for its true value. This implies that the effective market size, which below we find to be larger than the village, would be even larger if we accounted for credit responses in the pass-through analysis.

We solve for n using again the two values of the treatment-control differences in effective prices. We obtain the following results: *i*) with $\lambda = 14$, n , is 12.7 [7.9,17.1]; *ii*) with $\lambda = 149$, n is 13.75 [90% C.I. is 7.6,20.1]. These estimates of n imply that, according to the model, traders behave as if the number of their competitors were substantially larger (sixty to eighty percent) than just the number of traders operating in the same village, 7.8. The difference is consistent with the idea that some farmers have the option to sell outside the village.⁵⁷

nature of the price data. High pass-through is also consistent with evidence that the impact of international price shocks is large enough to affect local economic activities, conflicts, and long-term health outcomes (see, e.g., Adhvaryu et al., 2013; Brückner and Ciccone, 2010; Bazzi and Blattman, 2014; Adhvaryu et al., 2014). It is however important to note that other studies find low level of pass-through in agricultural supply chains. Examples include Fafchamps and Hill (2008) and Mitra et al. (2013).

⁵⁶In practice, we estimate Γ and n jointly.

⁵⁷The option to sell to other traders shapes competition, not the actual number of traders actually purchasing from each farmer. Anecdotally, we find that sales outside the village may be indirect. For instance, a farmer may give her product to a local aggregator who then makes sales outside the village. Consistent with the idea that villages do not necessarily match the relevant market size, 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. Alternative specifications that use the level, the inverse hyperbolic sine transformation, or dummies for the number of treated traders and of study traders give similar results (since our baseline survey has information on the number of traders but not on market shares, we cannot run spillover regressions by baseline treatment market shares). However, as a caveat, we note that these

6 Model Validation and Analysis

In this section, we first validate our framework choice by presenting additional evidence in support of our model and by showing that alternative models are inconsistent with the data. We also discuss how to extend our estimation method to account for heterogeneity along several margins. Second, we use the model to quantify the impact of the experiment on prices, quantities and farmer revenues (relative to a counterfactual without the experiment). Third, we look at the impact of counterfactual experiments that target a different share of traders. In doing so, we highlight the importance of specifically estimating the differentiation rate, rather than just the pass-through rate.

6.1 Model Validation

Robustness

The estimate of the difference between prices paid by subsidized and unsubsidized is a test of the null hypothesis of firm homogeneity (i.e. zero differentiation) in a wide range of models of imperfect competition. However, estimating market structure parameters does obviously require committing to a specific model, which in our case is admittedly a very stylized one (we discussed the key assumptions in Section 3.3). We provide evidence in support of our model choice. We proceed in three steps. First, since the model is overidentified, we can use different moments to recover the same parameters of interest.⁵⁸ We derive theoretical expressions for the *percent* difference in prices and quantities between treatment and control traders and match them to their empirical counterparts (0.007 and 1.88, respectively). Using again the pass-through rate ρ as an additional moment, we recover the differentiation rate, Γ , the market size, n , and the intercept parameter, α (this latter is only identified up to a monetary unit parameter). The estimates are very similar to the previous ones. For instance, when using our estimates of $\lambda = 210$, we obtain $\Gamma = 0.181$ and $n = 13.82$, which are very close to our previous estimates.⁵⁹

The fact that the results are so close is not mechanical. The first, more parsimonious, approach uses *level* differences in prices between treatment and control; the second method uses percent differences (both in prices and quantities). There is no mechanical relationship

estimates suffer from low power. Results are available on request.

⁵⁸Appendix C includes the details of this alternative approach.

⁵⁹Results are similar when using $\lambda = 149$.

between the two sets of moments. Indeed, Appendix Figures C.1 and C.2 show that the estimates of the two parameters would be very different when using other arbitrary values for the treatment-control percent price difference in a neighborhood of the real ones (for given level differences).

Alternative Models

Second, we discuss alternative models. While our baseline model is Cournot, we also implemented the steps described above using Bertrand competition (while retaining other assumptions on producers and buyers). The procedure delivers unrealistic values (a value of Γ larger than one and a market size n between 1 and 2). 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).⁶⁰ Quantity constraints (arising, for instance, from transport technologies) may be relevant in this setting.

Another candidate model could be collusion: the fact that treatment-control differences are small for prices and high for quantities may be consistent with treatment and control buyers forming a cartel to take advantage of the subsidy.⁶¹ However, 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. Collusion is also inconsistent with the large differential response of treatment traders in terms of advance payment provision and with the high pass-through rate of industry-wide price shocks presented in Section 5.1. In addition, Figure 3 does not show any stark or sudden decline in quantities purchased by control traders during the intervention period (which we would expect if traders were colluding to take advantage of the subsidy).⁶²

⁶⁰We note however that our model assumes independence of traders' choices across villages. This allows for quantity constraints at the trader-village level, but it rules out quantity constraints on the total volume of trader purchases (i.e., across villages).

⁶¹See Brooks et al. (2016) for a novel strategy to test collusive behavior.

⁶²Obviously, one caveat to this claim is that it relies on simple comparisons of the weeks right before or after the intervention and thus it does not account for potential changes in counterfactual quantities. The key point is however that control purchases do not fall drastically, which we would expect with collusion. In Section 6.2, we use our estimates and calculate that control purchases fall by around one-quarter.

We also consider standard models 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.⁶³

Heterogeneity

Finally, we discuss the role of heterogeneity in the estimation procedure. Appendix A.1 provides details. First, the estimation can accommodate trader heterogeneity in resale prices and, thus, heterogeneous trader size at baseline. Second, the estimation can accommodate heterogeneous levels of differentiation. Specifically, we can account for multiple symmetric locations with differentiation varying with the distance across locations (i.e. traders located further away from each other are more differentiated). In this case, the estimation procedure with pooled data recovers the differentiation rate across traders competing in the same location and a weighted “effective market size”, where the weight of each competitor is decreasing in its distance (and, thus, in its differentiation).

Arbitrary heterogeneity across locations in market size and differentiation is harder to accommodate. However, in principle, one could estimate Γ and n 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.⁶⁴

6.2 The Impact of the Experiment on Prices, Quantities, and Revenues

Having gained confidence in the choice of the framework, we now use the model to quantify the impact of the experiment on prices, quantities, and farmer revenues, relative to a counterfactual scenario without the experiment’s subsidies. We also study the impact of

⁶³Trader 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.

⁶⁴For 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.

counterfactual experiments that subsidize different share of traders. Appendix D provides details.

Prices

First, through the model, we can study how the subsidy affected (effective) prices traders pay to farmers. The derivative of the prices of treatment and control traders are pinned down by Γ , n , and μ (the share of treatment traders in the market). Figure 5 shows, for the estimated values of the competition parameters (i.e., $\Gamma = .176$, $n = 13.75$), 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, $\mu \in (0, 1)$.

At $\mu = 0.2$ —the share of treated traders in the market once including the non-study traders—treatment (control) traders raise their effective prices by 0.30 (0.15) per unit of subsidy and thus by 46 (23) Leones overall in response to the experimental subsidy of 150 Leones.⁶⁵ This exercise confirms that that the observed experimental difference between treatment and control, 23 Leones, is the result of a (partial) price war induced by giving the subsidy to a share of traders.

Quantities

During the experimental period, treatment traders purchase substantially larger amounts than control ones (+188%). Here, we aim to understand which share of this increase comes from market stealing vs. increases in aggregate supply. The model, through the *direct* supply function (which again depends on Γ and n), provides a mapping from the above price impacts to the quantity impacts (see Appendix D).

First, we find that *aggregate* quantity increases by an upper bound of 0.9%.⁶⁶ We can compare this result to the increase in aggregate quantity that would occur if all of the quantity results came from increases in aggregate supply (as opposed to market stealing), 38%. This suggests indeed that the difference in quantities between treatment and control traders during the experiment arises almost entirely from market stealing. This is consistent with the observation that, as the experiment was implemented at harvest time, farmers had

⁶⁵In setting $\mu = 0.2$, we are assuming that non-study traders are equal to control ones. The results in this section are based on the value $\lambda = 210$.

⁶⁶The lower bound is, trivially, zero.

limited options to increase their supply in response to the price changes.⁶⁷

Second, we can then assess the impact of the experiment on quantities purchased by treatment and control traders. Using the upper (lower) bound on the increase in aggregate quantity, we find that treatment traders increased their purchases by 111% (109%), relative to the counterfactual without the experiment, and control traders decreased their purchases by 27% (27.5%).

Farmer Revenues

The percent increase in farmer revenues is pinned down by the percent changes in prices and quantities for treatment and control, weighted by their market shares. Using again the lower (upper) bound on changes in aggregate supply, we find that farmer revenues increase by 1.1% (1.2%)

6.3 Counterfactual Experiments with Different Treatment Shares

In the intervention described in this paper, only 20% of the traders received the subsidy. We now compare the impact of alternative subsidy programs that target different treatment shares. One way to think about the subsidy is as an income transfer program to farmers. In particular, we are interested in comparing different subsidy interventions along their “return on investment:” the ratio between benefits in terms of incremental farmer revenues, and their cost (the total subsidy value). The continuous blue curve in Figure 6 presents the results using our estimates $\Gamma = .18, n = 13.75$, under the assumption of no impact of the experiment on aggregate quantity.⁶⁸ We note several points. First, once shutting down the aggregate quantity impact, the return is obviously always less than one: the subsidy value is passed only imperfectly to farmer revenues. Second, the return is increasing in the share of treated traders: the additional benefits in terms of farmer revenues when increasing μ are larger than the extra costs. Third, the return on investment is quite flat in the share of treated traders μ . In particular, under the estimated market structure parameters, the return when subsidizing 20% of the traders is about two-thirds of the return when targeting all the traders.

These counterfactual results may be useful to inform subsidy policies. For instance, a government may wish to transfer income directly to farmers, but to do so may be costly,

⁶⁷While farmers could not expand production dramatically, they could still increase output by harvesting marginal fruits or reducing processing losses.

⁶⁸As discussed above, the upper bound for the increase in quantity was 0.9%.

given the absence of financial infrastructure. If the cost of physically going out to transfer cash to farmers is more than the loss incurred by passing the subsidy through traders on to farmers, the government may find that paying a subsidy to traders is a (second-best) efficient way to transfer income to farmers. A helpful benchmark is the unconditional cash transfer (UCT) program of *Give Directly*, studied by Haushofer and Shapiro (2016). This program, which relies on mobile money technology to disburse payments, achieves a ratio of recipient benefits to costs of 94.7% in Kenya and 93.2% in Uganda.⁶⁹ These programs have a higher ratio of recipient benefits to total costs than trader subsidies for any level of μ . However, in Sierra Leone, only 4.5% of the population over 15 years has a mobile money account (Demirgüç-Kunt et al., 2015), and so such returns may not be obtainable.

These counterfactuals are also useful to highlight the importance of estimating the two market structure parameters separately. For a given level of the pass-through rate, different pairs (Γ, n) lead to different returns on investment when the share of treated traders is less than one. For instance, in Figure 6, the dashed orange curve captures the return for $\Gamma = 0.75$ and $n = 43$. While the orange and blue curve take, by construction, the same value for a “pass-through experiment” that targets all traders (i.e., $\mu = 1$), the orange line is above the blue line for interventions that only target a subset of traders. In particular, at $\mu = 0.2$, the return is about fifteen percentage points higher than before.

7 Conclusion

Most of the evidence on the competitiveness of agricultural markets in Sub-Saharan Africa relies on price-cost margins and on price dispersion analysis. In this paper, we have developed an experimental approach that combines randomized trader subsidies and a standard model of imperfect competition to recover key market structure parameters. We show that, in the presence of competitive spillovers, the treatment-control price difference can inform an intuitive test of the degree of differentiation across firms: only if firms are differentiated, there can be systematic price differences between subsidized and non-subsidized firms. In our study setting, we recover a degree of differentiation among traders of 0.1-0.2 on a scale from 0 (i.e., perfectly homogeneous traders) to 1 (i.e., each trader is a local monopolist). We hope this approach will prove useful for the many cases in which market-level randomization

⁶⁹These numbers are based on costs reported for 2014 by *Give Directly*, and calculated as $ROI = 1 - (\text{Cost of identifying recipients, transferring money, and following up}) / (\text{Total Cost} - \text{Overhead})$, to be comparable to our ROI measure, which includes only direct costs of the subsidy, and not overhead.

with appropriate clusters is not feasible for economic, logistical, or budgetary reasons.

In addition, by using a quasi-experimental estimate of the pass-through rate as an additional moment, we can infer the number of traders that compete for supply, rather than assuming it. We recover an effective market size of 13-14 traders —about 75% larger than the number of traders operating in a village (7.8). This finding may have implications for the design of cluster randomized controlled trials that attempt to generate variation at the market level: the relevant number of economic actors may be greater than those physically observed in a location. While this analysis relies on a specific model, several robustness exercises provide support for the model choice.

Overall, our results suggest a competitive intermediary sector, at least for those traders who purchase from farmers. This suggests that low differentiation and a high level of competitiveness can be found in settings where traders provide credit to their suppliers. This may be useful to consider the relevance of our findings for other settings: financial service provision along the supply chain does not necessarily induce a strong supplier base segmentation across firms. When considering external validity, we must also note that lower levels of the value chains (e.g. wholesalers, exporters), may be substantially less competitive than traders.⁷⁰ Understanding the degree of competitiveness at these other levels is an important question for future research.

⁷⁰Graphical evidence in Figure 2 suggests that, while wholesaler prices respond to changes in the international price, pass-through is substantially more incomplete at these lower levels of the supply chain.

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.
- Adhvaryu, Achyuta, James Fenske, and Anant Nyshadham.** 2014. “Early life circumstance and adult mental health.” *University of Oxford Department of Economics Discuss Paper Series ISSN 1471–0498*.
- Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham.** 2013. “Booms, busts, and household enterprise: Evidence from coffee farmers in tanzania.” Technical report, Working Paper.
- Aghion, Philippe, Jing Cai, Mathias Dewatripont, Luosha Du, Ann Harrison, and Patrick Legros.** 2015. “Industrial policy and competition.” *American Economic Journal: Macroeconomics*, 7(4): 1–32.
- 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.
- Allen, Treb.** 2014. “Information frictions in trade.” *Econometrica*, 82(6): 2041–2083.
- Antràs, Pol, and Arnaud Costinot.** 2011. “Intermediated trade.” *The Quarterly Journal of Economics*, 126(3): 1319–1374.
- Armstrong, Mark, and John Vickers.** 2015. “Which demand systems can be generated by discrete choice?” *Journal of Economic Theory*, 158 293–307.
- Atkin, David, and Dave Donaldson.** 2015. “Who’s Getting Globalized? The Size and Implications of Intra-national Trade Costs.” Technical report, National Bureau of Economic Research.
- Attanasio, Orazio, and Elena Pastorino.** 2015. “Nonlinear pricing in village economies.” Technical report, National Bureau of Economic Research.
- Baird, Sarah, J Aislinn Bohren, Craig McIntosh, and Berk Ozler.** 2014. “Designing experiments to measure spillover effects.”
- Bardhan, Pranab K.** 1980. “Interlocking Factor Markets and Agrarian Development: A Review of Issues.” *Oxford Economic Papers*, 32(1): 82–98, URL: <http://ideas.repec.org/a/oup/oxecpp/v32y1980i1p82-98.html>.
- Bardhan, Pranab, Dilip Mookherjee, and Masatoshi Tsumagari.** 2013. “Middlemen margins and globalization.” *American Economic Journal: Microeconomics*, 5(4): 81–119.
- Bardhan, Pranab, and Christopher Udry.** 1999. *Development microeconomics.*: OUP Oxford.

- Bazzi, Samuel, and Christopher Blattman.** 2014. “Economic shocks and conflict: Evidence from commodity prices.” *American Economic Journal: Macroeconomics*, 6(4): 1–38.
- Bell, Clive.** 1988. “Credit markets and interlinked transactions.” In *Handbook of Development Economics*. eds. by Hollis Chenery, and T.N. Srinivasan, 1 of Handbook of Development Economics: Elsevier, , Chap. 16 763–830, URL: <http://ideas.repec.org/h/eee/devchp/1-16.html>.
- Bergquist, Lauren Falcao.** 2017. “Pass-Through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya.” Technical report, Mimeo, UC Berkeley.
- 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.** 2017. “Strategic default in the international coffee market.”
- Breza, Emily, and Cynthia Kinnan.** 2016. “Measuring the Equilibrium Impacts of Credit: Evidence from the Indian Microfinance Crisis.” *Manuscript, Northwestern University*.
- Brooks, Wyatt J, Joseph P Kaboski, and Yao Amber Li.** 2016. “Growth Policy, Agglomeration, and (the Lack of) Competition.” Technical report, National Bureau of Economic Research.
- Brückner, Markus, and Antonio Ciccone.** 2010. “International commodity prices, growth and the outbreak of civil war in Sub-Saharan Africa.” *The Economic Journal*, 120(544): 519–534.
- 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.** 2014. “Selling low and buying high: An arbitrage puzzle in Kenyan villages.”
- Busso, Matias, and Sebastian Galiani.** 2014. “The causal effect of competition on prices and quality: evidence from a field experiment.” Technical report, National Bureau of Economic Research.
- 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*.
- 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.** 2016. “Firm and Market Response to Saving Constraints: Evidence from the Kenya Dairy Industry.” Technical report.
- Casaburi, Lorenzo, and Jack Willis.** 2016. “Time vs. State in Insurance: Experimental Evidence from Contract Farming in Kenya.”
- Chau, Nancy H, Hideaki Goto, and Ravi Kanbur.** 2016. “Middlemen, fair traders, and poverty.” *The Journal of Economic Inequality*, 14(1): 81–108.
- 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.” *The Quarterly Journal of Economics*, 128(2): 531–580.
- Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran.** 2015. “The price effects of cash versus in-kind transfers.” Technical report, National Bureau of Economic Research.
- 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*.
- Demirgüç-Kunt, Asli, Leora F Klapper, Dorothe Singer, and Peter Van Oudheusden.** 2015. “The global finindex database 2014: Measuring financial inclusion around the world.”
- Dillon, Brian, and Chesley Dambro.** 2016. “How competitive are food crop markets in sub-Saharan Africa?” *Available at SSRN 2752748*.
- 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.” *The American Economic Review*, 67(3): 297–308.
- 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.** 2017. “Credit Rationing and Pass-Through in Supply Chains: Theory and Evidence from Bangladesh.”
- 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.
- 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.** 2017. “Relationships, risk and rents.”
- Gilbert, Christopher L et al.** 2009. “Cocoa market liberalization in retrospect.” *Review of business and economics*, 54(3): 294–312.
- Haushofer, Johannes, and Jeremy Shapiro.** 2016. “The short-term impact of unconditional cash transfers to the poor: Experimental Evidence from Kenya.” *The Quarterly Journal of Economics*, 131(4): 1973–2042.
- 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.
- Jaffe, Sonia, and E Glen Weyl.** 2010. “Linear demand systems are inconsistent with discrete choice.” *The BE Journal of Theoretical Economics*, 10(1): .
- Jensen, Robert.** 2007. “The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector.” *The quarterly journal of economics*, 122(3): 879–924.
- Jensen, Robert, and Nolan Miller.** 2016. “Information, Demand and the Growth of Firms: Evidence from a Natural Experiment in India.”

- Kreps, David M, and Jose A Scheinkman.** 1983. “Quantity precommitment and Bertrand competition yield Cournot outcomes.” *The Bell Journal of Economics* 326–337.
- Lalive, Rafael, Camille Landais, and Josef Zweimüller.** 2015. “Market externalities of large unemployment insurance extension programs.” *The American Economic Review*, 105(12): 3564–3596.
- Macchiavello, Rocco, and Ameet Morjaria.** 2015a. “Competition and Relational Contracts: Evidence from Rwanda’s Mills.” *Unpublished*.
- Macchiavello, Rocco, and Ameet Morjaria.** 2015b. “The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports.” *American Economic Review*, forthcoming.
- Maitra, Pushkar, Sandip Mitra, Dilip Mookherjee, Alberto Motta, and Sujata Visaria.** 2014. “Financing Smallholder Agriculture: An Experiment with Agent-Intermediated Microloans in India.”
- McKenzie, David, and Susana Puerto.** 2017. “Growing Markets through Business Training for Female Entrepreneurs.”
- Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria.** 2013. “Asymmetric information and middleman margins: An experiment with west bengal potato farmers.” Technical report.
- 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.** 2016. “Experimentation at Scale.” Technical report, Working Paper.
- 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.
- Rotemberg, Martin.** 2017. “Equilibrium Effects of Firm Subsidies.” Technical report, Mimeo, Harvard University.
- Rubin, Donald B.** 1974. “Estimating causal effects of treatments in randomized and non-randomized 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.
- Startz, Meredith.** 2016. “The value of face-to-face: Search and contracting problems in Nigerian trade.”
- 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

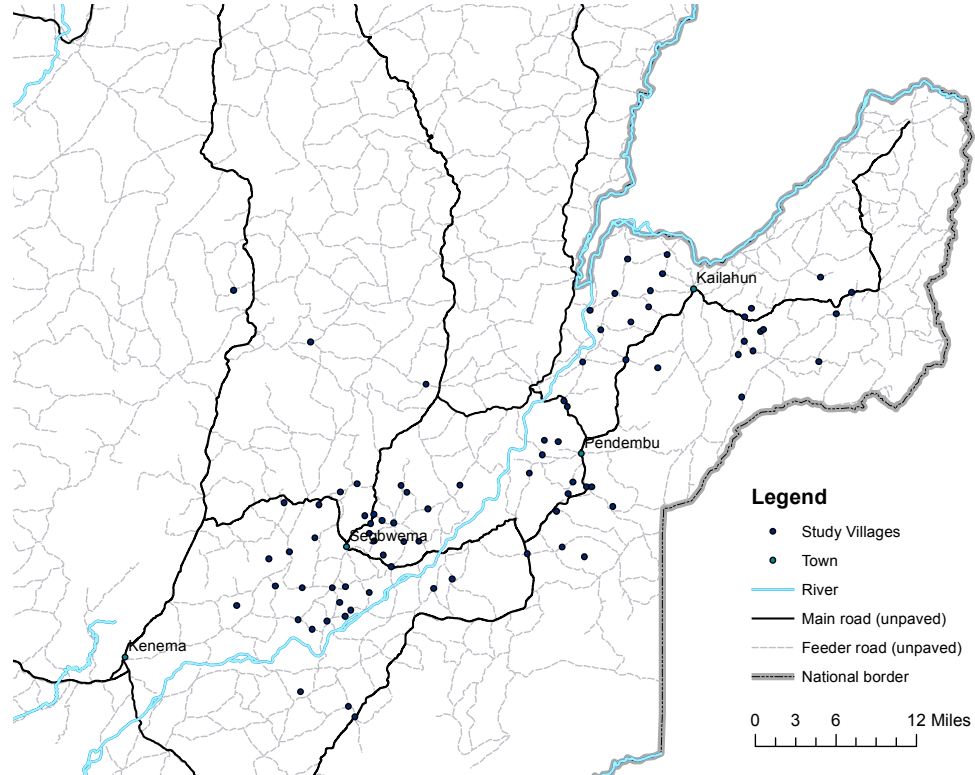




Figure 2: Cocoa Prices




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. Wholesaler and trader prices 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: Estimating λ : Treatment-Control Differences by Chiefdom, Prices vs. Advance Payments




Notes: The scatter reports the 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 -271.

Figure 5: Counterfactual Experiments: Trader Prices



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, $\mu = 0.2$. At this value of the treatment share, control (treatment) traders increase their prices by 0.15 (0.30) per unit of subsidy, respectively, relative to a scenario without the experiment. For $\mu \rightarrow 1$, the response of treatment traders tends to the pass-through rate, 0.92.

Figure 6: Counterfactual Experiments: Return on Investment



Notes: The graph reports the return to investment (y-axis) for experiments that target a share μ (x-axis) of traders with a unitary per-unit subsidy. The return to investment is defined as the ratio between the additional farmer revenues generated by the intervention and the cost of the subsidy. We conduct the simulations assuming no response in aggregate supply. Section 6.2 describes the procedure to recover these values. The continuous curve describes the results using the market structure parameters estimated in the paper ($\Gamma = .18, n = 14$). The dashed curve describes the results using an alternative pair ($\Gamma = .75, n = 43$) that gives the same pass-through rate, $\rho = .92$, than the pair of values estimated in the paper.

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)	345 (694)	339 (762)	6.2 (96.5)

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 (56 traders). 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 Prices

	(1)	(2)	(3)	(4)	(5)
Bonus $\in \{0,1\}$	-32.52 (47.16)	-5.47 (14.95)	-5.92 (16.99)	-12.87 (13.21)	-6.86 (15.41)
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 Leone 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	0.14*** (0.03)	0.14*** (0.02)	0.13*** (0.03)	0.14*** (0.02)
Control Group Mean	0.11	0.11	0.11	0.11
Trader Controls		X		X
Village Controls			X	X
Observations	1837	1825	1837	1825

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 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 (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	537.12***	527.72***
	(54.0)	(54.2)
Control Group Mean	282.5	282.5
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 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. 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: 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: 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 6: The Value of Advance Payments: Heterogeneity in Treatment-Control Differences

	(1)	(2)	(3)
Treat* Estimated Treatment Effect on Credit	-341.79	-300.13	-209.87
<i>p-values from bootstrapped t-stats</i>	[.10]	[.15]	[.07]
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 $\pi_a^{\tilde{p}}$ from equation 13. P-values in brackets are derived from pairs cluster bootstrap-t at the randomization pair level using 1,000 replications. 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.

Table 7: Pass-Through from Wholesaler to Trader Prices

	OLS	IV		
	(1)	(2)	(3)	(4)
Wholesaler Price	0.89*** (0.02)	0.92*** (0.01)	0.93*** (0.01)	0.90*** (0.09)
Control Group Mean	3007	3007	3007	3007
Kleibergen-Paap First Stage F-stat		1408.2	471.6	5.5
Trader FE	X	X	X	X
Month FE			X	X
Village FE				X
Observations	1254	1254	1254	1254

Notes: The table reports the pass-through from wholesaler prices (i.e. paid to traders) to trader prices (i.e., paid to farmers). Both are measured in Leones per pound. An observation is a shipment delivered by the trader to a wholesaler before or during the intervention. In Columns (2)-(4), wholesaler prices are instrumented with the front-month prices for liquid cocoa futures, obtained by the *Intercontinental Exchange*. Standard errors are clustered by day. ***p<0.01, **p<0.05, *p<0.1.

Appendix

A Theory Appendix

A.1 Trader Heterogeneity

The baseline model presented in Section 3 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.⁷¹ 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), (6), and (7) 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 $\Delta q = \frac{s}{\beta(1+\Gamma)}$). 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} .⁷²

A.2 Non-study Traders

As discussed above, the model presented in Section 3 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

⁷¹This is equivalent to varying producer costs in an oligopoly model.

⁷²The 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 $\tilde{n} = \frac{n}{M} \frac{1-\kappa^M}{1-\kappa}$.

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 $p_i = \alpha + \beta q_i + \gamma \sum_{j \neq i} q_j$.⁷³ 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 3 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).

B 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 CAO-BISCO/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.

⁷³That is, we assume a common degree of differentiation across study and non-study traders.

C Recovering Competition Parameters: An Alternative Approach

This Appendix presents details about the alternative approach to recover Γ and n that we presented in Section 6.1. Our goal is to identify alternative (quasi-)experimental 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

Our main estimation approach relies on two moments: the *level* difference in treatment and control prices (Equation 5) and the pass-through rate of changes in wholesaler prices (Equation 7). In this section, we show how the key parameters Γ and n , and 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 price at which traders purchase is Le. 2,987, 91% 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.⁷⁴

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 7)—, 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 price of 2,987. The average likelihood of advance provision for control traders is 0.11. Therefore, given $\lambda = 209$, the average control effective price is 3,010. This implies that the percent price difference between treatment and control traders during the experiment is 0.7%. The average quantity purchased by control traders is 282.5 kilograms. Thus the percent different between treatment and control traders is 188%.


⁷⁴Results are quite stable when using other values of v , spanning between 2,987 (the average trader price) and 3,260 (the average wholesaler price).

Solving the equation system with these values for $\% \Delta p$ and $\% \Delta q$, we obtain the following estimates for the three parameters of interests: $\Gamma = 0.181$, $n = 13.8$, and $\alpha = 2,015$. 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 since one 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 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 4.3.

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 of the steps to assess the impact of the experiment on prices, quantities, and farmer revenues (Section 6.2).

Setup

We use the superscript 0 to refer to the pre-intervention period and 1 to refer to the intervention period. At baseline, traders are homogeneous and pay p^0 .⁷⁵ Each trader thus faces the direct supply. Thus their direct supply function is $q_i^0 = a + bp_i^0 - c \sum_{i \neq j} p_j^0$.⁷⁶ Symmetry implies $q^0 = a + (b - c(n - 1))p^0$. Aggregate supply is thus $Q^0 = nq^0 = n[a + (b - c(n - 1))p^0]$.

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 $dp_T = p_T^1 - p^0 = 46$ and $dp_C = p_C^1 - p^0 = 23$.⁷⁷ Using a baseline price of 2,964 (mean of the effective price for control traders during the experiment minus dp_C), we obtain $p_T^1/p^0 = 1.015$ 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 1.5% (0.8%).

Impact on Quantities

Given dp_T and dp_C , we can write: $q_T^1 = a + b(p^0 + dp_T) - c((\mu n - 1)(p^0 + dp_T) + (1 - \mu)n(p^0 + dp_C))$. With some algebra, we obtain $q_T^1 = q^0 + \underbrace{(bdp_T - c((\mu n - 1)dp_T + (1 - \mu)ndp_C))}_{dq_T}$. Similarly, for the control group,

$q_C^1 = q^0 + \underbrace{(bdp_C - c(\mu ndp_T + ((1 - \mu)n - 1)dp_C))}_{dq_C}$. The aggregate quantity during the experimental period,

Q^1 , is then $Q^1 = n(\mu q_T^1 + (1 - \mu)q_C^1) = n(\mu(q^0 + dq_T) + (1 - \mu)(q^0 + dq_C)) = Q_0 + \underbrace{n(\mu dq_T + (1 - \mu)dq_C)}_{dQ^1}$. The

increase in aggregate quantity induced by the experiment can be written as: $dQ^1 \equiv Q^1 - Q^0 = n(b - c(n - 1))(\mu dp_T + (1 - \mu)dp_C)$. In turn, the *percent* impact is: $\frac{dQ^1}{Q^0} = \frac{n}{n} \frac{(b - c(n - 1))(\mu dp_T + (1 - \mu)dp_C)}{a + (b - c(n - 1))p^0} = \frac{\mu dp_T + (1 - \mu)dp_C}{\frac{a}{b - c(n - 1)} + p^0}$.

Since we do not estimate β and γ separately, but only their ratio, we cannot quantify $\frac{dQ^1}{Q^0}$. However, assuming

⁷⁵Thus, throughout the exercise, we assume non-study traders and study traders are homogeneous before the experiment.

⁷⁶The direct supply function is $q_i = a + bp_i - c \sum_{j \neq i} p_j$, with $a \equiv \frac{\alpha}{\beta + \gamma(n - 1)}$, $b \equiv \frac{\beta + \gamma(n - 2)}{(\beta + \gamma(n - 1))(\beta - \gamma)}$, $c \equiv \frac{\gamma}{(\beta + \gamma(n - 1))(\beta - \gamma)}$.

⁷⁷This is, by construction, consistent with our estimate of the difference in (effective) prices between treatment and control traders.

$a \geq 0$ (which holds in our estimates) and noticing that $b - c(n - 1) > 0$ (since $\beta > \gamma$), then $\frac{\mu dp_T + (1 - \mu) dp_C}{p^0}$ is an *upper bound* on $\frac{dQ^1}{Q^0}$.

From the percent impact on aggregate quantity, we can now compute the impact for treatment and control quantities. We do this in a four steps

1. Aggregate quantity is $Q^1 = n(\mu q_T^1 + (1 - \mu)q_C^1) = nq_0(1 + \frac{dQ^1}{Q^0})$
2. We define the ratio of treatment to control quantities during the experimental period: $R \equiv q_T^1/q_C^1$. Then,

$$Q^1 = n(\mu R q_C^1 + (1 - \mu)q_C^1)$$

3. We can now solve for q_C^1 (relative to q^0):

$$\frac{q_C^1}{q^0} = \frac{1 + \frac{dQ^1}{Q^0}}{\mu R + (1 - \mu)}$$

We note that we can measure the ratio R in the data

4. Finally, we can easily derive $\frac{q_T^1}{q^0} = R \frac{q_C^1}{q^0} = R \frac{1 + \frac{dQ^1}{Q^0}}{\mu R + (1 - \mu)}$

In our experiment, $\mu = .2, 1 - \mu = .8, dp_T = 46, dp_C = 23. p^0 = 2,964$. Therefore we compute the upper bound on the percent change in aggregate quantity as

$$\frac{\widehat{dQ^1}}{Q^0} = \frac{27}{\frac{a}{b - c(n - 1)} + 2987} \leq \frac{27.6}{2987} = 0.009$$

In response to the experiment, aggregate quantity raises by *at most* 0.9%. Also, we have a trivial lower bound, which is zero (i.e., the aggregate quantity is constant). As a benchmark, we can quantify the increase in aggregate quantity we would observe without any market stealing (i.e., the control quantities were unchanged). In this case aggregate quantity would go up by $\mu * 188\% + (1 - \mu) * 0$, that is by 38%.⁷⁸ This suggests that most of the difference between treatment and control in quantity purchases comes indeed from market stealing. Finally, given the upper bound of 0.009, we can compute that, relative to a world without experiment, controls reduce their purchases by 27% and treatment increase their purchaes by 111%. At the lower bound of 0, controls reduce their purchases by 27.5% and treatment traders increase their purchaes by 109%.

Impact on Farmer Revenues

In the pre-experiment period, farmer revenues are simply $r^0 = p^0 Q^0 = p^0 n q^0$. In the experimental period, these become $r_1 = n(\mu p_T^1 q_T^1 + (1 - \mu) p_C^1 q_C^1)$. Therefore, the ratio between these two values is a function of the quantities we derived above:

$$\begin{aligned} \frac{r_1}{r_0} &= \frac{\mu p_T^1 q_T^1 + (1 - \mu) p_C^1 q_C^1}{p^0 n q^0} \\ &= \mu \frac{p_T^1}{p^0} \frac{q_T^1}{q^0} + (1 - \mu) \frac{p_C^1}{p^0} \frac{q_C^1}{q^0} \end{aligned}$$

⁷⁸188% is the percent difference in treatment and control quantities during the experiment.

In our experiment, taking the upper bound $dQ^1/Q^0 = 0.009$, we obtain:

$$\frac{r_1}{r_0} = \mu \frac{p_T^1}{p^0} \frac{q_T^1}{q^0} + (1 - \mu) \frac{p_C^1}{p^0} \frac{q_C^1}{q^0} = .2 * 1.015 * 2.11 + .8 * 1.007 * .73 = 1.02$$

At the lower bound of no change in aggregate quantity, $\frac{r_1}{r_0} = 1.01$.

Returns on Investment and Counterfactual Experiments

Finally, we consider the return on investment (ROI) on experiments 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. The former is $r^1 - r^0 = \underbrace{\left(\mu \frac{p_T^1}{p^0} \frac{q_T^1}{q^0} + (1 - \mu) \frac{p_C^1}{p^0} \frac{q_C^1}{q^0} \right) - 1}_{r^1/r^0} \underbrace{np^0 q^0}_{r^0}$. The cost of the intervention is

$C = \mu nsq_T^1$. The derivations in the paper consider the case in which the experiments do not induce an increase in aggregate quantity. We focus first on variations in μ given our estimates of (Γ, n) and then consider returns for alternative values of these parameters. Section 6.2 presents the results.