Competition in Agricultural Markets: An Experimental Approach*

Lorenzo Casaburi **Tristan Reed**

April 2017

Abstract

This paper presents an experimental approach to measure competition in agricultural markets, based on the random allocation of subsidies to competing traders. We compare prices of subsidized and unsubsidized crop traders to recover the key market structure parameter in a standard model of imperfect competition. By combining the experimental results with quasi-experimental estimates of the pass-through rate, we also estimate market size, or the effective number of traders competing for farmers' supply. In the context of the Sierra Leone cocoa industry, our results point to a competitive agricultural trading sector and suggest that the market size is substantially larger than the village. The methodology developed in this paper uses purely *individual-level* treatment to shed light on market structure. This approach may be useful for the many cases in which market-level randomization is not feasible.

*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, Pascaline Dupas, Fred Finan, Matthew Gentzkow, Robert Gibbons, Rachel Glennerster, Oliver Hart, Asim Khwaja, Michael Kremer, Rocco Macchiavello, Ted Miguel, Ben Olken, Dina Pomeranz, Orie Shelef, Tavneet Suri, Chris Udry, Jack Willis, Josef Zweimüller and workshop participants at CSAE Oxford, 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 a recent survey of the empirical literature on competition in Sub-Saharan Africa, Dillon and Dambro (2016) note however that the evidence is "remarkably thin." In the limited cases where conclusions can be drawn, the authors see a high degree of competition. But on the other hand, governments and international institutions often emphasize the monopsonistic power of traders (see, e.g., Kedir, 2003; Gresser and Tickell, 2002). Previous studies of competition in these markets 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), 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). 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. By comparing prices that subsidized and unsubsidized traders pay to farmers, we can recover the rate of differentiation between traders, the key market structure parameter in a standard model of oligopsonistic competition. In intuitive terms, this parameter measures the extent to which prices paid by competitors affect the quantity supplied to a trader, relative to the trader's own price sensitivity.[1] The main 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. Further, we combine our experimental results with quasi-experimental estimates of the pass-through rate of common price shocks across all traders to estimate market size, or the number of traders effectively competing for farmer supply.

An ideal experiment to measure competition would identify separate markets and randomize the subsidy treatment at the market level. For instance, in two influential papers, Weyl and Fabinger (2013) and Atkin and Donaldson (2015) show that, in a general model

[1] More precisely, since our framework features quantity-setting traders, the differentiation rate depends on the ratio between the slope of a trader's inverse supply to her competitors' quantity and the slope to her own quantity.

of symmetric imperfect competition, the pass-through rate—i.e., the difference in prices between two markets with an infinitesimal difference in producer costs—is the key object to recover the degree of competition in the market. However, this approach is often unfeasible as researchers face two common constraints in the design of market-level RCTs. First, choosing an appropriate definition of truly independent markets can easily imply that the unit of observation is so large that there are too few units over which to randomize (or that randomization across markets would be very expensive). Second, a firm may operate in several locations and only partially overlap with the operation areas of other firms. In this case, it is hard to identify market boundaries that partition the space and thus to define appropriate randomization clusters. For instance, in our setting, traders operate in multiple villages and they often face different competitors in each of these villages. Motivated by these widespread challenges, this paper proposes a method that uses purely *individual*-level randomization at the trader level to study market structure.

We use a standard framework of oligopsonistic competition among traders to guide 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"). We model the experiment by introducing a unit subsidy on the wholesaler price to a subset of traders. Two parameters shape competition: the degree of differentiation among traders, which captures a preference of farmers to sell to specific traders, and the market size, defined as the number of traders to which farmers have the option to sell. We characterize equilibrium trader prices and quantities for subsidized and non-subsidized traders, and, crucially, the differences in equilibrium outcomes between the two groups of traders.

The model highlights the strategic interaction between subsidized and unsubsidized traders that compete in the same market: 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, since randomization still ensures that subsidies are uncorrelated with firm characteristics, these differences depend only on the experimental subsidy and they can be used to recover market parameters. Specifically, the difference between prices paid by (randomly) subsidized and unsubsidized traders depends on the differentiation rate: in a market with perfectly homogeneous traders, there is one price, thus there can be no systematic difference between the prices paid by the two groups. If traders are differentiated, however, different prices can coexist and unsubsidized traders can pay a lower price.[2] Similarly, the extent to which a subsidized trader can steal market shares from unsubsidized competitors also depends on differentiation. We use these intuitive relations to recover the differentiation parameter.

In the experiment, treatment traders—about 20% of the traders operating in the study region—receive from the wholesaler a per-unit subsidy worth about 5% of the average baseline trader price.[3] We find the following results. During the experimental period, treatment traders pay similar prices to farmers than control traders. They are also more likely to provide advance payments (i.e., liquidity services) to farmers that they report as regular suppliers at baseline (+14 percentage points, a 117% increase). Field interviews suggest that, in this market, the willingness to offer advance payments for cocoa is an important source of differentiation between traders.[4] Using two different strategies to measure the value of advance payments—baseline correlations and treatment heterogeneity—we compute the difference between treatment and control traders in the "effective price," a price akin to the net present value of the transaction. This difference 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 are between 0.1 to 0.2, on a 0-1 scale (where 0 implies complete homogeneity and 1 implies that each trader operates as a monopsonist).[5]

This differentiation rate, which we recover from the experimental results, summarizes the level of competition in the model *for a given number of traders*. However, we acknowledge that a geographically defined village in which farmers live may not be the appropriate definition of a market, and so the number of traders cannot be known precisely ex-ante, for instance by counting traders in a village. Farmers may have the option to sell to traders that operate primarily in other villages, or they may sell only to traders that come to their farm

[2]With complete differentiation, the difference in prices between treatment and control traders equals the pass-through rate of a monopsonist.

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

[4]This observation is consistent with the significant literature on interlinked transactions in agriculture, which emphasizes the role of buyers in providing inputs and credit services. Bardhan (1980), Bell (1988), and Bardhan and Udry (1999) summarize a large body of theory that relates land, labor, output, and credit markets. Blouin et al. (2013), Casaburi and Willis (2016), Casaburi and Macchiavello (2016), Ghani and Reed (2014), and Macchiavello and Morjaria (2015a) provide primarily empirical contributions. Our paper emphasizes the relationship between interlinked transactions and price transmission. Other sources of differentiation may include search costs, repeated relationships, and trader quantity constraints.

[5]In 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.

gate.[6] We therefore treat the number of traders competing for farmer supply as an additional parameter to recover.[7] We estimate the number of traders by combining the experimental results with quasi-experimental estimates of the pass-through rate from wholesaler to trader prices (i.e. of a common cost shock affecting all traders). Specifically, we estimate the pass-through rate using plausibly exogenous variation in world cocoa prices and find that its value is 0.92. Given this result, the "effective market size" we recover is 14 traders. This implies that the number of traders competing for a farmer's cocoa is about 80% larger than the average number of traders observed operating in a village (7.8). In other words, traders behave as if the number of their competitors was substantially larger than just the number of traders operating in the same village.[8] This finding points to the difficulty of defining market size ex-ante, as would be done implicitly by partitioning space in a market-level randomization.

Overall, consistent with the survey by Dillon and Dambro (2016), the estimates of the differentiation rate and of the market size suggest that the market is quite competitive. Indeed, in response to the subsidy, treatment traders purchase substantially more cocoa than control farmers (+188%). With the model, we can 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. Further, we find that the difference in quantities purchased by treatment and control traders arises primarily from market stealing, as opposed to increases in aggregate supply, which we can bound at 0.9%. Specifically, the experiment increased the quantity purchased by treatment traders by about 110% and decreased that of control traders by about 27%. The lack of change in aggregate quantity is consistent with the fact that the experiment began halfway in the harvest season and thus farmers had limited options to increase their output volume.[9] By combining the price and quantity responses for control and treatment traders, we then compute that total farmer revenues increased by 1 to 2%.

Finally, we quantify the impact on farmers of counterfactual subsidy programs treating

[6]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. Therefore, collecting information from farmers may not be sufficient to trace the relevant market size.

[7]In terms of the model, the market size is the number of traders competing for supply of the homogeneous producers of a village.

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

[9]These options include harvesting marginal fruits (e.g., at the top of the trees) or reducing processing losses.

different shares of traders, given the value of the competition parameters we recovered. We define the "return on investment" of these subsidies as the ratio of incremental farmer revenue in response to the subsidy to the value of subsidy paid to the traders. Our experiment, which treated around 20% of the traders in the study region, achieved two-thirds of the return of an intervention providing the same unit subsidy to *all* traders. In addition, for a given level of pass-through, the welfare implication of a subsidy targeting only a subset of traders varies with the differentiation rate and the market size. This emphasizes the importance of estimating both these competition parameters and not just their combination, which the pass-through rate captures.

The methodology developed in this paper does require committing to a specific model of imperfect competition, which in our case is admittedly a very stylized one. While the insights for the estimation are likely to generalize to other frameworks, the equations for the equilibrium solutions depend on the functional forms of the model. We provide additional evidence to mitigate concerns related to the specific choice of the model. First, we are able to use different moments (i.e., the *percent* differences between treatment and control traders in prices and quantities) to recover the two market structure parameters. The parameter estimates we obtain are very close to those estimated using the first set of moments (we also show that this result is not mechanical). Second, we present evidence suggesting that alternative models such as Bertrand, monopsonistic competition, and collusion among traders are not consistent with the data.

To the best of our knowledge, this is the first experiment that randomized any treatment at the trader level in agricultural markets (Dillon and Dambro, 2016) and also the first to use experimental subsidies to measure competition. Besides contributing to the literature on the competitiveness of agricultural markets, the paper relates to a growing body of work that seeks to estimate the equilibrium effects of (quasi-) experiments using market-level randomization.[10] Several recent papers are particularly related to our work. Rotemberg (2014) identifies the spillover effects on larger firms of a subsidy available only to small firms in India, using variation in exposure to eligible competitors. Busso and Galiani (2014) studies the effect of increased competition on prices and quality of goods, using randomized entry of new retailers in the Dominican Republic. Mitra et al. (2013) uses the results of a farmer price information intervention in West Bengal to test among alternative models of farmer-

[10]Other 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. (2016).

intermediary trades. Falcao Bergquist (2016) randomizes subsidies to maize consumers and traders in Kenya, in order to infer market structure parameters of the Atkin and Donaldson (2015) framework.[11] Relative to this body of work, our paper provides a proof of concept for how individual-level randomization can shed light on market parameters. We hope this approach will prove useful for the many cases where market-level randomization is not feasible for the reasons described above.

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 results from the randomized controlled trial and the quasi-experimental estimates of pass-through. Section 5 combines the theoretical and empirical results to recover market structure parameters and to quantify the impact of the experiment. Section 6 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 country's largest export crop by value, according to the UN COMTRADE database. 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).[12] Farmers sell to

[11] Other examples include Crépon et al. (2013), Cunha et al. (2015), Lalive et al. (2015), Baird et al. (2014), Mitra et al. (2013), Hildebrandt et al. (2015), Burke (2014), Mobarak and Rosenzweig (2014), Muralidharan et al. (2016).

[12] Sierra 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). Though Sierra Leone does have an official marketing board, the organization has been defunct since the 1990s, and the government is responsible for a negligible share of purchases. A potential explanation for the lack of vertical integration in the market in the absence of a strong marketing board are the stringent legal restrictions on the transaction of land discussed in Acemoglu et al. (2014). These, along with weak legal institutions more broadly, would make vertical integration of the supply chain difficult, if not impossible.

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.[13] The analysis in this paper focuses on grade A cocoa, the grade targeted by the experimental subsidy, unless otherwise specified.

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.[14] These wholesalers collect cocoa in their warehouses using a network of trader with whom they have exclusive relations (i.e. a trader almost always delivers cocoa to only one wholesaler). Wholesalers then sell the cocoa 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

[13] Appendix B provides details.

[14] These towns are now quite remote, accessible only by unpaved roads that can become impassible in the rainy season. During the colonial period, however, Pendembu was a prosperous trading town and the final stop on the Sierra Leone Railroad, which was dismantled and sold by the government of Siaka Stevens in the 1974. The decline in the country's cocoa industry since then can be observed at the massive abandoned produce warehouse where the end of the tracks once lay. Exporters we visited joked with some cynicism that the cocoa stocks of the largest wholesalers in Pendembu could not come close to filling it.

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.[15]

2.3 Data Collection and Summary Statistics

Over the course of the experiment we collected a variety of data from traders.[16] 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.[17] On average, based on the trader survey, there are 7.8 traders in a village but only 3.2 study traders, suggesting that about 60% of the traders in the market are non-study traders (i.e. working with other wholesalers). In Section 5.1, 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. 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.[18] 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

[15] Of the 84 traders identified by wholesalers, four were outliers with respect to baseline quantity relative to other traders (within the same wholesaler), and could not be matched to other traders in our randomization strategy. Thus, the final sample selected for randomization was 80 traders.

[16] One important caveat to our empirical analysis is that, due to funding constraints, all the data are self-reported by traders when they visit the wholesaler warehouses.

[17] Figure 1 presents a map of the study setting.

[18] Interviews 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).

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 our research team measured and documented the quality of their shipment. We collected data starting 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. In the shipment data, we collected the price per pound paid to farmers, and the name of the village where the cocoa in the shipment 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. In addition, in order to study the impact of trader treatment on provision of advance payments, 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).

3 A Simple Model of Buyers' Imperfect Competition

This section presents a standard 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:[19]

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

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

We follow an extensive literature in industrial organization that uses linear demand to model imperfectly competitive industries (see, e.g., Vives, 2001).[20] We adapt the standard model of differentiated producers to our setting, which features imperfect competition among buyers. In the model, differentiation allows in a reduced form for potential unobserved aspects of a specific buyer-seller relationship that may give the buyer market power.

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.[21]

Buyers

The profit of buyer i in a village is given by

$$\pi_i = q_i(v_i - p_i), \quad (2)$$

where q_i is the quantity purchased, v_i is the (net) resale price, and p_i is the price paid to producers.[22]

We assume buyers are ex-ante symmetric in the resale price v.[23] The experiment introduces a subsidy, s, for a share μ of the buyers. 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).

---

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.

[20] As 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.

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

[22] A 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.

[23] Appendix A.1 relaxes this assumption.

We assume *Cournot oligopsonistic competition*: each buyer sets quantities strategically, taking into account competitors' choices. In Section 5.2, we provide several arguments 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.

Equilibrium Quantities and Prices

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

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

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

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

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).

Treatment-Control Differences in Quantities and Prices

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

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

and

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

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.[24] On the other hand, Δq is *decreasing* in Γ (for given β): if traders are homogeneous, the treatment traders can expand 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.

Pass-Through

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

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

which is decreasing in Γ and increasing in n.

3.3 Recovering the Competition Parameters

By matching the theoretical expressions in Equations (5) and (7) to their empirical analogues (which we derive in Section 4), we can recover the competition parameters of the model. Specifically, for a given value of the subsidy, Equation (5)—i.e., the experimental estimate of the treatment-control difference in prices—enables us to recover the degree of differentiation among traders, Γ. For a given the number of competing traders, n, the experimental results are thus sufficient to characterize competition, since this depends only on Γ.

[24] If $\Gamma = 1$, 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 $\Delta p = s/2$.

However, we treat market size as an additional parameter to recover. This deserves further explanation. The number of traders physically operating in a village would seem a natural starting point to define market size.[25] However, farmers may also sell outside their villages. For instance, some farmers may bring their cocoa to other locations directly or through other farmers and certain market places may serve multiple villages. Alternatively, farmers may only sell to traders that come to the farm gate. Therefore, the relevant market size for a farmer may be either larger or smaller than the number of traders who make purchase in the village. For this reason, we choose to treat the market size as an unknown rather than assuming its value. Specifically, given the estimate of Γ, Equation (7)—i.e., the quasi-experimental estimates of the pass-through rate—allows us to recover n.[26]

3.4 Discussion of the Assumptions

The stylized model presented in this section makes strong assumptions. We discuss some of the key ones. First, it 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.[27] Second, we rely on specific functional forms. For instance, the linear supply function of the model 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).[28] 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.[29] 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

[25]For instance, most villages have a meeting point where sales occur.

[26]We note that this estimation procedure does not use the treatment-control differences (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 take as a parameter to estimate. Section 5.2 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.

[27]We discuss this topic further in Section 4.2.

[28]The general model of symmetric imperfect competition presented in Atkin and Donaldson (2015) is not readily applicable to the asymmetry introduced by the randomized subsidy (also, refer to Weyl and Fabinger (2013) for a discussion of the complications that arise when modeling asymmetric firms.)

[29]For instance, Attanasio and Pastorino (2015) present evidence of nonlinear pricing in rural Mexican villages and propose a model of price discrimination to account for this nonlinearity.

instance, if the subsidy leads to entry or exit.[30]

Because of these restrictive assumptions, it is important to provide evidence in support of the choice of the model. We do this in Section 5.2. 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. In addition, Appendix A.1 shows that the estimation framework is robust to allowing for baseline heterogeneity among traders. We leave to future research to generalize the framework by relaxing the other assumptions discussed above.

Finally, we note that differentiation among traders can originate from a range of reasons, such as search costs, repeated relationships, quantity constraints, and transport costs. Section 5 suggests 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. Therefore, Γ should be interpreted as a "reduced form" parameter, which captures several of these forces.

4 Experimental and Quasi-Experimental Results

This section presents the experimental and quasi-experimental results that will be used to recover competition parameters in the next section. First, it reports treatment-control differences in trader prices, advance payment provision, and quantities purchased during the intervention period. Second, it quantifies the value of advance payments and thus the treatment-control difference in *effective prices*, akin to the net present value of the payment. Finally, it presents estimates of the pass-through rate of industry-wide shocks in wholesaler prices to trader prices, based on exogenous changes in the world price of cocoa.

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.

[30]In another example, the subsidy could relax liquidity constraints of the 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).

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. We have data on trader prices at the shipment level.[31] These values include payments that are (potentially) made at different times (e.g., before and after harvesting). 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.[32]

Figure 2 displays the price results graphically. It shows weekly averages for: i) world prices (right y axis);[33] ii) wholesaler prices; iii) trader prices paid by treatment traders; 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:

$$\tilde{p}_{sizt} = \eta_z + \eta_t + \pi^{\tilde{p}}\text{Treat}_i + \epsilon_{sit}, \tag{8}$$

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).[34]

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

[31]The analysis presented in this section uses the modal price paid for cocoa in the shipment, as reported by the trader. We also verify that results are similar in a specification in which the outcome is an alternative measure of price taken by dividing the shipment total expenditure by its weight.

[32]We recover the treatment-control differences for effective prices in Section 4.2.

[33]Specifically, 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.

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

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, when prices were lower.[35]

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, one at a time and then combined, controls referring to the trader and to the village where the majority of cocoa in the shipment originated.[36] Trader controls include the number of suppliers the trader buys 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 include baseline share of suppliers begin given credit in the village, number of other bonus traders and number of study traders, miles to nearest town, number of clients across all traders, and fixed effects for village chiefdoms. Columns (3)-(5) in Table 2 show that the coefficient of interest is quite stable when including these controls. This suggests that the selection concerns described above cannot drive the results.[37]

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. In Section 5.3, 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).

[35]As discussed in Section 2, six traders never visited the warehouses. Effectively, by including pair fixed effects in the regression, we drop their pairs form the analysis.

[36]In the intervention period, 80 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.

[37]In results not presented, we also tested for effects on the prices of B and C grade cocoa. Though we found no significant treatment-control difference in the price of grade C cocoa, we did 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.

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:

$$AdvancePayment_{fiz} = \eta_z + \pi_1^a \text{Treat}_i + \nu_{fi} \tag{9}$$

An observation is a farmer listed as a regular supplier in the trader baseline. $AdvancePayment_{fiz}$ is an indicator of whether farmer f was given advance payments by 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 percentage points. 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.

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, quantities purchased are balanced between treatment and control in the three weeks before the intervention. Second, throughout the intervention, treatment traders purchase substantially higher volumes than control ones. Third, the graph shows an increase in the total quantity purchased by study-traders 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 the time the intervention ended.

Table 4 presents regression results from the following regression models:

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

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.[38] The results are robust when including trader controls in Column (2). Overall, this is a large impact of the treatment.[39] Given that farmers had limited opportunities to increase production by the time the intervention started, it seems likely that the results are primarily driven by market-stealing effects. The analysis in Section 5.3 reaches a similar conclusion using the model.

4.2 The Value of Advanced Payments

This section quantifies the value buyers provide to farmers through the provision of advance payments. It then combines the previous results on treatment-control differences in prices and advances to estimate the treatment-control difference in the *effective price*, which, as discussed, is a measure akin to net present value and that 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 moments in order to estimate the competition parameters.

It is important to emphasize that our framework does not model the trader's choice to pass value to the farmer through advance payments versus simply raising the trader price. Our model does not explain why, during the experiment, treatment traders are more likely to pay in advance than control ones, but do not pay a higher price at harvest time.[40] As discussed in Section 3.4, the static framework obviously does not capture the repeated game nature of these 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

[38]We 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.)

[39]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.

[40]We 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.

and it may feature multiple equilibria, thus substantially complicating the estimation.

While we do not explicitly model the trader choice to provide advances, it is nevertheless important to compute the value of these advance payments. We therefore propose a simple method of introducing the additional value of advance payments into our framework. Ideally, 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. Therefore, we define the effective price paid by buyer i as it follows:

$$p_i = \tilde{p}_i + \lambda \cdot ShareAdvances_i, \tag{11}$$

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. 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).[41] We estimate the value of λ twice separately using two different sources of variation: cross-sectional (at baseline) and experimental.

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.[42] 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).[43] 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

[41] Since 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.

[42] In this section, 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 5). 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).

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

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. In other words, we estimate how much less a trader who increases her advance payments must adjust her prices.

For this purpose, we modify Equation (9) to allow for heterogeneity in the treatment-control differences across villages and trader characteristics:

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

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 difference in advance payment provision between treatment and controls 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:

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

If total payments and advance payments are substitutes (i.e., $\tilde{\lambda} > 0$), then $\pi_a^{\tilde{p}} < 0$.[44]

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

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

geographic units of local legal and political administration, and, as discussed in Acemoglu et al. (2014) vary in contract enforcement and other institutions.[45] 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 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.

Overall, both approaches (cross-sectional baseline correlations and treatment heterogeneity) suggest that farmers value advance payments and are willing to take a lower price when offered an advance. One must note that neither of these two approaches provides *causal* estimates of λ. For instance, villages with higher shares of advance provision may be systematically different and have some feature (not captured by our controls) that lowers prices.[46] It is however reassuring that both approaches lead to positive and significant value of λ, even if they exploit 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.[47] In addition, assuming that the loan covers the entire purchase and considering

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

[46]In 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.

[47]We 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 is used in the first approach, captures whether the farmer had received advances in the twelve

an average loan duration of 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.[48]

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: $\hat{\pi}^{p} \equiv \hat{\pi}^{\tilde{p}} + \hat{\lambda} \cdot \hat{\pi}^{a}$. Based on Tables 2 and 3, we specify $\hat{\pi}^{\tilde{p}} = -6.9$ and $\hat{\pi}^{a} = .14$. The values of λ are 150 when using the baseline correlations and 210 when using the treatment heterogeneity (specifically, 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 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). Importantly, as discussed above, this difference could be the result of price competition between traders, and it does not necessarily 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 $\hat{\pi}^{\tilde{p}}, \hat{\lambda}, \hat{\pi}^{a}$. We obtain [-10.93,37.53] and we can reject the null that the different in effective price is zero at p=0.23. For the second approach based on treatment heterogeneity, we instead use bootstrap, clustering resampling at the randomization pair level. The confidence interval is [-31.11, 61.33]. The latter result highlights the lower precision of the estimates from the second approach, as we have to account for the fact that $\pi_{a}^{\tilde{p}}$ is a generated regressor.

4.3 Pass-Through of Industry-Wide Price Shocks

In this section, we study how trader prices respond to common changes in wholesaler prices. As preliminary evidence of high responsiveness, Figure 2, discussed in Section 4.1.1,

months before the baseline, a longer time horizon that the one of the experiment (two months). Thus, the two advance payment dummies may capture different intensities of advance payments. An ideal data collection would have gathered transaction-level information on advance and harvest payments.

[48] According 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. We also note that our data and empirical strategy do not allow us to compute transaction-level interest rates and thus we cannot quantify the impact of the subsidy on the terms of the loan.

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. Throughout the table, standard errors are clustered by date. The coefficient estimate is 0.91. This suggests a high-level of pass-through from traders to farmers. Note that these results are not directly comparable with those in studies that focus on pass-through of *border* prices to final producers/consumers, such as in Atkin and Donaldson (2015).[49]

The change in wholesaler prices may be correlated with local supply shocks. We implement two strategies to mitigate this concern. First, we instrument wholesaler prices with the international price of cocoa, as measured by the *Intercontinental Exchange*. Given that Sierra Leone has a very small share of the global production and that other major producers harvest at different times of the year, 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 estimates increases to 0.92. Second, starting from Column (3), we add month fixed effect and, thus, only rely on variation within the same month. The coefficient is stable. Finally, Columns (4) and (5) show that the coefficient is robust to the gradual inclusion of trader and village fixed effects (though the first stage becomes substantially weaker when adding village fixed effects).[50]

5 Recovering Competition Parameters

This section combines the model of Section 3 and the experimental and quasi-experimental results of Section 4 to recover competition parameters of the model. First, we discuss the relation (and the discrepancies) between the field experiment and the model. Second, we report the estimates for Γ and n. Third, we present additional evidence in support of the framework and discuss alternative models. Fourth, we use the results to discuss the impact of the experiment (and of counterfactual experiments) on prices, quantities, and farmer margins.

[49]Studies of pass-through to producer prices specifically in Sub-Saharan Africa include, among others, Fafchamps and Hill (2008), Adhvaryu et al. (2013), and Dillon and Barrett (2015).

[50]Graphical 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.

5.1 The "Theory Experiment" and the "Field Experiment"

The simple model presented in Section 3 derives prices and quantities for subsidized and unsubsidized traders in a static equilibrium. The field experiment setting obviously presents some deviations from this stylized environment. Here, we discuss such discrepancies and discuss how they may affect the estimation. 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 competition we estimate with our intervention would have been higher 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 of contestable farmers and thus we may have observed a larger price response.[51]

Second, the experiment only ran until the end of the harvest season. Traders may have behaved differently in a multi-season trial. For instance, treatment traders may have increased their prices less than they would have done in a longer experiment, in an attempt to capture the short run value of the subsidy. For this reason, 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 4.3 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 for 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 implicit 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, who 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 however is robust to the presence of such traders. Specifically, in Appendix A, we discuss a model where only a share σ of traders is included in the study, and thus study

[51] As discussed 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.

treatments are a share $\sigma\mu$ of traders. Non-study traders have a resale price that possibly differs from that of the study traders, v. We show that the equilibrium theory treatment-control differences, Δp and Δq (Equations 5 and 6), and the pass-through rate ρ (Equation 7) are unchanged.[52] Therefore, the estimation approach is robust to the inclusion of non-study traders.

Finally, we discuss two additional assumptions necessary to reconcile the theory to the available data. First, we assume that advance payments to *regular* suppliers—those on which we have data—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. Second, 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 then assume that the price pass-through we measured in Section 4.3 is equal to the pass-through of the effective price (i.e., we assume advance payments do not respond to transitory world price shocks). This may not be a particularly realistic assumption since, in our experiment, advance payment provision does respond to changes in the resale price (however, in the experimental case, the price change was guaranteed to hold for the entire season, as opposed to the case of transitory shocks). In light of this consideration, our estimate of ρ (0.92) may be considered a lower bound (and thus our estimate of n would also be a lower bound for the real value).

5.2 Recovering Γ and n

Results

Following the approach described in Section 3.3, we use the empirical results of the previous section to recover the competition parameters of the model. By matching the theoretical expressions in Equations (5) and (7) to their empirical analogues, we obtain a system of two equations in two unknowns, Γ and n:

$$\begin{cases} \Delta p & \equiv & \frac{s\Gamma}{1+\Gamma} & = & \hat{\pi}_1^p + \hat{\tilde{\hat{\lambda}}} \cdot \hat{\pi}_1{}^a \\ \rho & \equiv & 1 - \frac{1}{1+\Gamma+n(1-\Gamma)} & = & \hat{\rho} \end{cases} \tag{14}$$

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

We solve the above system using the two values of the treatment-control differences in effective prices that we derived in Section 4.2, $\lambda = 149$, from the baseline cross-section, and $\lambda = 210$, from the heterogeneity in treatment-control differences. We obtain the following results for point estimates and 90% confidence intervals: *i)* the estimated value of the differentiation rate Γ is 0.10 [-.10,.29] with the first method and 0.176 [-.18,.63] with the second method; *ii)* the estimate of the market size parameter, n, is 12.7 [7.9,17.1] with the first method and 13.75 [7.6,20.1] with the second method.

We can always reject complete differentiation and never reject homogeneity (i.e., $\Gamma = 0$).[53] The estimated value of n is sixty to eighty percent higher than the average number of traders (study and non-study) observed in a village, 7.8. The difference is consistent with the idea that some farmers may sell, directly or indirectly, outside the village. The difference stresses the importance of treating the market size as an unknown variable.[54] Overall, the results suggest fairly competitive markets. This is in line with the review by Dillon and Dambro (2016) and is also consistent with the fact that, at baseline, the ratio between the price traders pay to the farmers and the price traders receive from the wholesaler is on average 0.92. Therefore, 0.08 is an upper bound on the profit markdown buyers obtain in the transaction (i.e., which may be lower because of transport and other transaction costs).

Robustness, Alternative Models, and Heterogeneity

The approach we use to recover market structure parameters in this section does require committing to a specific model, which in our case is admittedly a very stylized one (as discussed in Section 3.4). We attempt to provide some evidence in support of our choice. We proceed in two steps. First, we propose an alternative approach that uses, in part, different moments to recover the same parameters of interest.[55] 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

[53] When we ignore that λ is a generated regressor, we can reject $\Gamma = 0$ at p=0.13.

[54] 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 estimates suffer from low power. Results are available on request.

[55] Appendix C includes the details of this alternative approach.

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.[56] 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 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).

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).[57] 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.[58] 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. We also note that collusion is inconsistent with the large differential response of treatment traders in terms of advance payment provision. It is also inconsistent with the high pass-through rate of industry-wide price shocks presented in Section 4.3 and with the low trader markdowns we discussed above. 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

[56] Results are similar when using $\lambda = 149$.

[57] We 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).

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

were colluding to take advantage of the subsidy).[59] 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.[60]

Finally, we discuss the role of heterogeneity in the estimation procedure (Appendix A.1 provides details). The estimation allows for trader heterogeneity in resale prices and, thus, for heterogeneous size at baseline. It also allows for 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 are less substitutable). In this case, the model can be used to recover the differentiation rate across traders competing in the same location and the "effective market size", where each competitor is weighted according to the inverse of its differentiation rate. 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.[61]

5.3 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 proceed in several steps. Appendix D provides details.

[59]Obviously, 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 5.3, we use our estimates and calculate that control purchases fall by around one-quarter.

[60]Trader 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.

[61]For 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.

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.[62] 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%.[63] 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 limited options to increase their supply in response to the price changes.[64]

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

[62] In setting $\mu = 0.2$, we are considering the case where non-study traders are equal to control ones. The results in this section are based on the value $\lambda = 209$.

[63] The lower bound is, trivially, zero.

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

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%)

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.[65] 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 to 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, 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

[65] As discussed above, the upper bound for the increase in quantity was 0.9%.

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.[66] 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.

6 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 obtain two key results. First, 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 monopsonist). The provision of advance payments is a candidate source of differentiation among traders. Second, 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 farmer supply. By treating market size as an unknown parameter, we acknowledge that the number of traders physically operating in a village may not be the relevant definition of the number of competitors. 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

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

observed in a location.

Overall, consistent with the survey by Dillon and Dambro (2016), our results suggest a competitive intermediary sector, at least for those traders who purchase from farmers. One important caveat is that lower levels of the value chains (e.g. wholesalers, exporters), may be substantially less competitive than traders. Understanding the degree of competitiveness at these other levels is an important question for future research. Methodologically, our analysis suggests that it is possible to learn about market-level features from *individual*-level randomization. We hope this approach will prove useful for the many cases in which a clustered experimental design is not feasible for economic, logistical, or budgetary reasons.

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, 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.

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.

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`.

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, Rocco Macchiavello et al. 2013. "Tropical lending: international prices, strategic default and credit constraints among coffee washing stations." *University of Warwick*.

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.

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 findex 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, M Helal Uddin et al. 2016. "Do Consumers Benefit from Supply Chain Intermediaries? Evidence from a Policy Experiment in Edible Oils Market in Bangladesh."Technical report, Boston University-Department of Economics.

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.

Falcao Bergquist, Lauren. 2016. "Pass-Through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya."Technical report, Mimeo, UC Berkeley.

Fold, Niels. 2005. "Global cocoa sourcing patterns." *Cross-continental Agro-food Chains: Structures, Actors and Dynamics in the Global Food System*, p. 223.

Ghani, T., and T Reed. 2014. "Competing for Relationships: Markets and Informal Institutions in Sierra Leone." *Unpublished.*

Gilbert, Christopher L et al. 2009. "Cocoa market liberalization in retrospect." *Review of business and economics*, 54(3): 294–312.

Gresser, Charis, and Sophia Tickell. 2002. *Mugged: Poverty in your coffee cup.*: Oxfam.

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): .

Kedir, Abbi Mamo. 2003. "Rural poverty report 2001: the challenge of ending rural poverty edited by the INTERNATIONAL FUND FOR AGRICULTURAL DEVELOPMENT (IFAD).(Oxford: Oxford University Press, 2001, pp. 266)." *Journal of International Development*, 15(5): , p. 667.

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."

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, 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. 2014. "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.

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

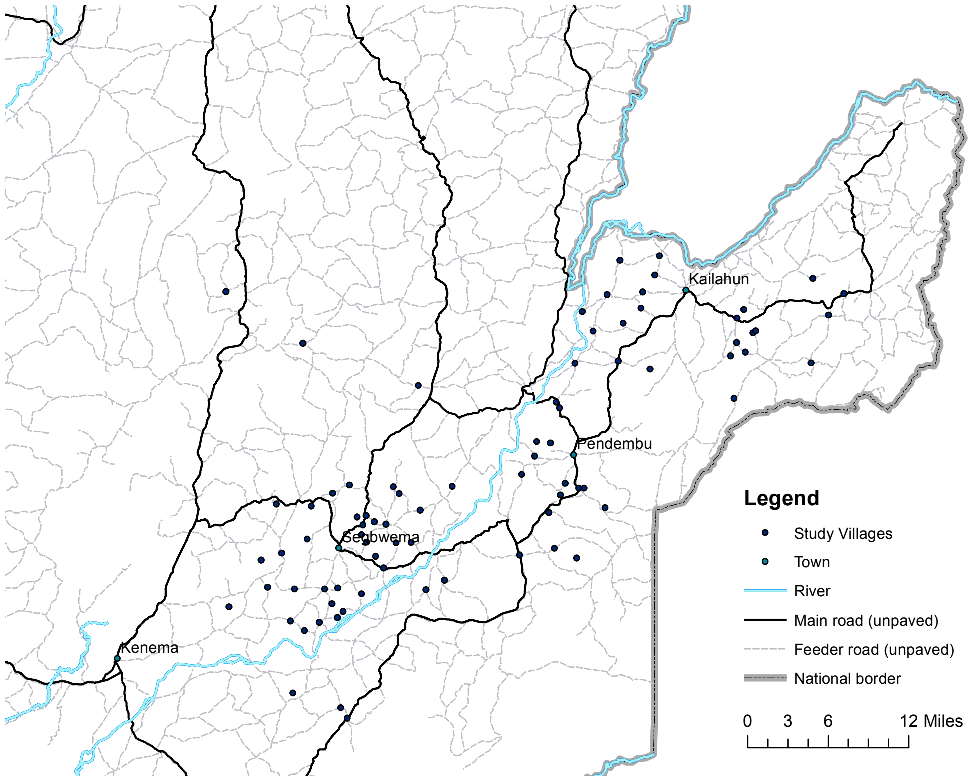

Figure 2: Cocoa Prices

*N*otes: 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

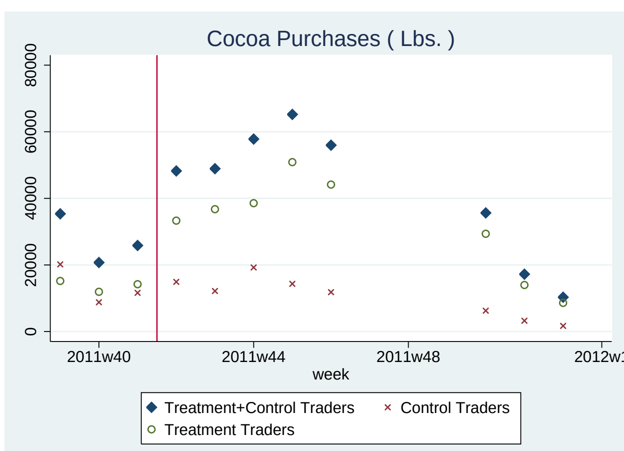

*N*otes: The figure shows the total amount of cocoa purchases by study traders (i.e., control and treatment traders), control traders, and treatment traders. The vertical line marks the beginning of the intervention period. Data collection was suspended for most of three weeks (w47-w49).

Figure 4: Estimating λ: Treatment-Control Differences by Chiefdom, Prices vs. Advance Payments

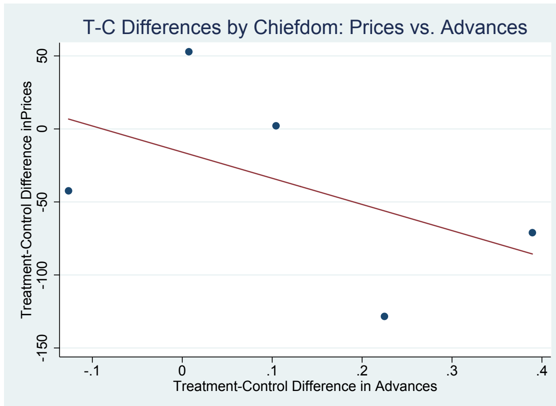

*N*otes: The scatter reports the the correlation across price and advance payments treatment-control differences, estimated separately across the five chiefdoms included in the study. The regression line has a slope of -271.

Figure 5: Counterfactual Experiments: Trader Prices

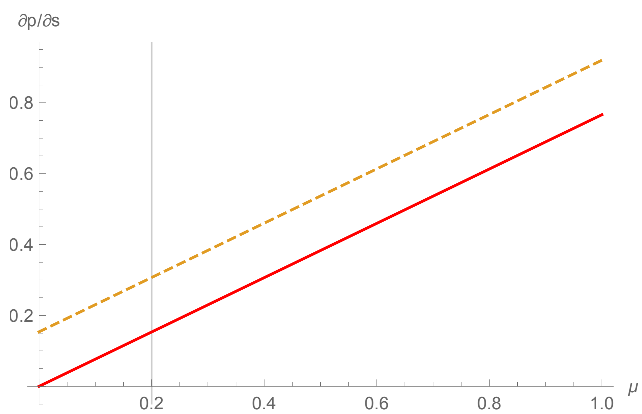

*N*otes: The graph shows the impact of counterfactual experiments on (effective) prices paid by control traders (continuous line) and treatment traders (dashed line). Specifically, it reports the increase in prices in response to a unit-subsidy as a function of the share of treated traders, μ. The vertical 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, relatively to a scenario without the experiment. For $\mu \to 1$, the response of treatment traders tends to the pass-through rate, 0.92.

Figure 6: Counterfactual Experiments: Return on Investment

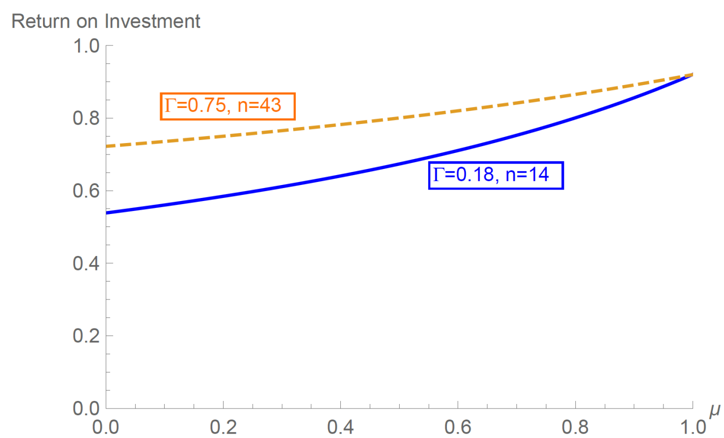

*N*otes: The graph reports the return to investment (y-axis) for experiments that target a share μ (x-axis) of traders with a unitary per-unit subsidy. The return to investment is defined as the ratio between the additional farmer revenues generated by the intervention and the cost of the subsidy. We conduct the simulations assuming no response in aggregate supply. Section 5.3 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
Panel A: Baseline Interview			
Self-estimate bags sold in 2011	20.0	18.6	1.5
	(28.3)	(18.5)	(2.23)
Age, years	38.2	36.9	1.4
	(8.2)	(10.2)	(1.91)
Years trading cocoa	8.1	8.9	-0.8
	(5.4)	(5.5)	(1.2)
Years selling to study wholesaler	5.7	7.3	-1.6
	(4.8)	(4.9)	(0.86)*
Cement or tile floor in house $\in \{0, 1\}$	0.53	0.63	-0.1
	(0.51)	(0.49)	(0.1)
Mobile phone owner $\in \{0, 1\}$	0.90	0.93	-0.03
	(0.30)	(0.27)	(0.06)
Access to storage facility $\in \{0, 1\}$	0.88	0.78	0.10
	(0.33)	(0.42)	(0.09)
Villages operating in	4.25	4.87	-0.62
	(1.64)	(2.02)	(0.39)
Number of suppliers per village	5.8	6.2	-0.35
	(3.3)	(3.6)	(0.84)
Share of suppliers given credit since March	0.72	0.68	0.04
	(0.32)	(0.28)	(0.05)
Panel B: Pre-treatment shipment data			
Price Paid to Farmer (shipment-level)	3,137	3,136	1.2
	(154)	(151)	(41.9)
Pounds sold during pre-treatment (weekly)	345	339	6.2
	(694)	(762)	(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 don 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: Price Responses to the Experiment

	(1)	(2)	(3)	(4)	(5)
Bonus $\in \{0,1\}$	-32.52	-5.47	-5.92	-12.87	-6.86
	(47.16)	(14.95)	(16.99)	(13.21)	(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: Advance Payments Responses to the Experiment

	(1)	(2)	(3)	(4)
Treatment Trader	0.14***	0.14***	0.13***	0.14***
	(0.03)	(0.02)	(0.03)	(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: Quantity Responses to the Experiment

	(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: Price vs. Share 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: Price vs. Advance Payments: Heterogeneity in Treatment-Control Differences

	(1)	(2)	(3)
Treat* Estimated Treament Effect on Credit	-341.79	-300.13	-209.87
p-values from boostrapped t-stats	[.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)	(5)
Wholesaler Price	0.91***	0.92***	0.94***	0.93***	0.90***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.09)
Control Group Mean	3007	3007	3007	3007	3007
Kleibergen-Paap First Stage F-stat		14024.4	2065.0	471.6	5.5
Month FE			X	X	X
Trader FE				X	X
Village FE					X
Observations	1254	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)-(5), 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.[67] 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: $\Delta p = \frac{s\Gamma}{1+\Gamma}$ 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 $\tilde{n}$.[68]

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 fundamentally different than the ones we include in the study. We present an extension of the model that

[67] This is equivalent to varying producer costs in an oligopoly model.

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

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 $p_i = \alpha + \beta q_i + \gamma \sum_{j \neq i} q_j$.[69] 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 CAOBISCO/ECA/FCC (2015) and, for the specific case of West Africa, David (2005). Other dimensions of quality affecting price on the international market are various fair-trade and environmental certifications. Such certification generally requires that beans can be verifiably traced to individual producers. In our market, there is not yet the infrastructure to do such tracing, and so this quality dimension does not apply.

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

[69] That 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 5.2. 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

The main approach described in Section 3.3 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:

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

and

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

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.[70]

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

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

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%.

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 approach we obtained in Section 5.2. 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 (Section 5.2), 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 $\frac{p_T - p_C}{p_C}$

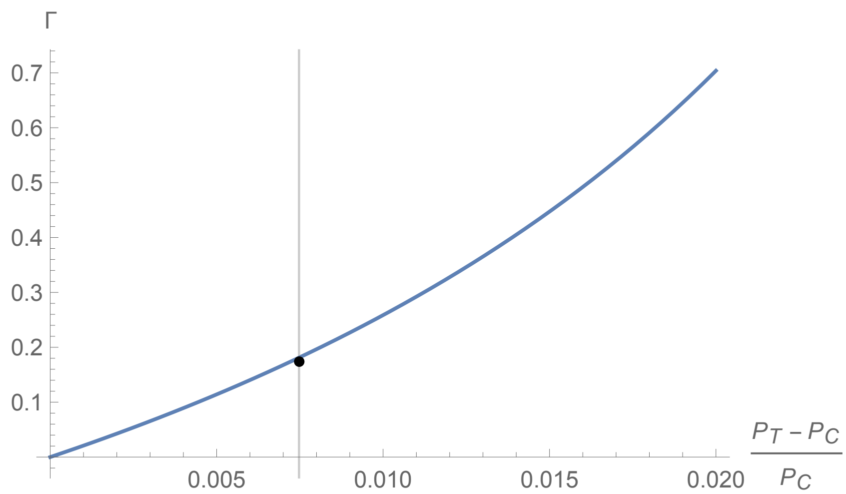

*N*otes: The graph reports sensitivity of the estimate of Γ obtained from the method described in Appendix C to the value of the percent treatment-control price difference. The dot represent the estimate from the main method presented in Section 3.3.

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

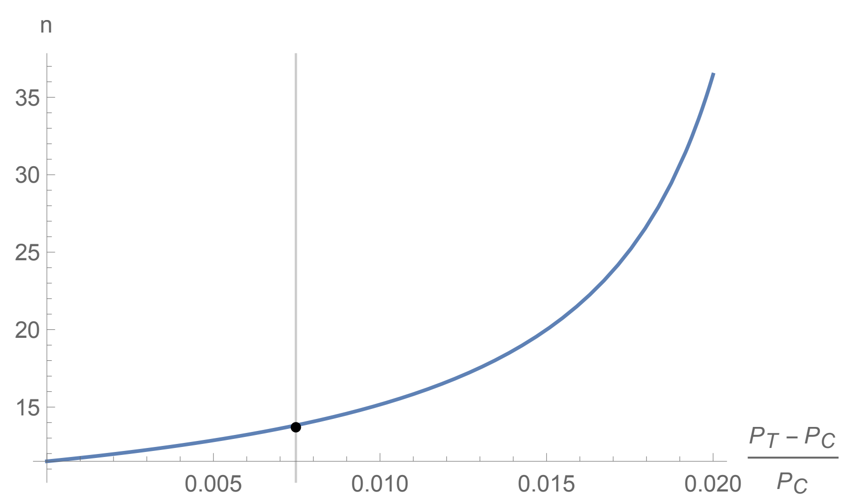

*N*otes: The graph reports sensitivity of the estimate of Γ obtained from the method described in Appendix C to the value of the percent treatment-control price difference. The dot represent the estimate from the main method presented in Section 3.3.

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 5.3).

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.[71] Each trader thus faces the direct supply. Thus their direct supply function is $q_i^0 = a + bp_i^0 - c\sum_{i\neq j} p_j^0$.[72] 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: $\frac{\partial p_T}{\partial s} = \frac{\Gamma - \frac{(\Gamma-1)\mu n}{\Gamma - \Gamma n + n + 1}}{\Gamma + 1}$, $\frac{\partial p_C}{\partial s} = \frac{(1-\Gamma)\mu n}{(\Gamma+1)(1+\Gamma+n(1-\Gamma))}$.

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 $dp_C = p_C^1 - p^0 = 23$.[73] 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: $q_T^1 = a + b(p^0 + dp_T) - c\Big((\mu n - 1)(p^0 + dp_T) + (1-\mu)n(p^0 + dp_C)\Big)$. With some algebra, we obtain $q_T^1 = q^0 + \underbrace{\Big(bdp_T - c((\mu n - 1)dp_T + (1-\mu)ndp_C)\Big)}_{dq_T}$. Similarly, for the control group, $q_C^1 = q^0 + \underbrace{\Big(bdp_C - c(\mu ndp_T + ((1-\mu)n - 1)dp_C)\Big)}_{dq_C}$. The aggregate quantity during the experimental period, Q^1, is then $Q^1 = n\big(\mu q_T^1 + (1-\mu)q_C^1\big) = n\big(\mu(q^0 + dq_T) + (1-\mu)(q^0 + dq_C)\big) = Q_0 + \underbrace{n\big(\mu dq_T + (1-\mu)dq_C\big)}_{dQ^1}$. The increase in aggregate quantity induced by the experiment can the be written as: $dQ^1 \equiv Q^1 - Q^0 = n\big(b - c(n-1)\big)\big(\mu dp_T + (1-\mu)dp_C\big)$. In turn, the *percent* impact is: $\frac{dQ^1}{Q^0} = \frac{n}{n}\frac{\big(b-c(n-1)\big)\big(\mu dp_T + (1-\mu)dp_C\big)}{a + (b - c(n-1))p^0} = \frac{\mu dp_T + (1-\mu)dp_C}{\frac{a}{b-c(n-1)} + p^0}$. Since we do not estimate β and γ separately, but only their ratio, we cannot quantify $\frac{dQ^1}{Q^0}$. However, assuming

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

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

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

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

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

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

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%.[74] 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 $r_1 = n\Big(\mu p_T^1 q_T^1 + (1-\mu)p_C^1 q_C^1\Big)$. Therefore, the ratio between these two values is a function of the quantities we derived above:

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

[74] 188% 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 $r^1 - r^0 = \left(\underbrace{\left(\mu \frac{p_T^1}{p^0}\frac{q_T^1}{q^0} + (1-\mu)\frac{p_C^1}{p^0}\frac{q_C^1}{q^0}\right)}_{r^1/r^0} - 1\right)\underbrace{np^0q^0}_{r^0}$. The cost of the intervention is $C = \mu n s q_T^1$. 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 5.3 presents the results.