Competition in Agricultural Markets: Experimental Evidence from Sierra Leone*

Lorenzo Casaburi **Tristan Reed**

August 28, 2016

Abstract

This paper proposes an experimental approach to study competition in agricultural markets. In a randomized controlled trial in Sierra Leone, a subset of cocoa traders received a per-unit subsidy on their purchases from farmers. A standard framework of imperfect competition among traders models the response of treatment and control prices and quantities to the subsidy. By matching the experimental treatment effects—as well as estimates of pass-through rate of wholesaler prices—to their theoretical counterparts, we find that: *i)* the market features a differentiation rate among traders of 0.25 (on a 0 to 1 scale) and that the provision of advance payments is a candidate source of differentiation; *ii)* the effective market size is twice as large as the village. The methodology developed in this paper suggests that, when combined with standard models, *individual-level* randomization can shed light on market structure.

JEL Classification: O13, Q14, F14

Keywords: Pass-through, interlinked transactions, intermediated trade, agricultural markets.

*Lorenzo Casaburi: casaburi@stanford.edu. Tristan Reed: treed@fas.harvard.edu. 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, Dave Donaldson, Pascaline Dupas, Fred Finan, Robert Gibbons, Rachel Glennerster, Gita Gopinath, Oliver Hart, Oleg Itshkoki, Michael Kremer, Ted Miguel, Esteban Rossi-Hansberg, Orie Shelef, Tavneet Suri, Chris Udry, and workshop participants at CSAE Oxford, Harvard, MIT, NBER, Paris School of Economics, UC Berkeley, UC San Diego and Stanford for helpful suggestions and comments. Derick Bowen, Grant Bridgman, Felix Kanu and Fatoma Momoh provided excellent research assistance. We gratefully acknowedge 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

Measuring the extent of competition in agricultural markets is essential for understanding how price incentives, taxes, or subsidies will propagate along the supply chain to the producers. In Sub-Saharan Africa, governments and international institutions often emphasize the monopsonistic power of traders (see, e.g., Kedir, 2003). A recent survey of the empirical literature by Dillon and Dambro (2016) concludes, however, that the bulk of evidence points to a competitive market, but with the caveat that evidence overall is limited. Existing studies have primarily relies on the analysis of price-cost margins at the trader level (Fafchamps et al., 2005 and Osborne, 2005), price dispersion (Fackler and Goodwin, 2001; Aker, 2010), and pass-through of international prices along the value chain (Fafchamps and Hill, 2008; Atkin and Donaldson, 2015).

In this paper, we instead propose an experimental approach based on a randomized controlled trial in Sierra Leone where a subset of traders received a subsidy per each unit of cocoa they purchase from the farmers, worth approximately 5% of the prices. We interpret the experimental results through the lens of a standard model of imperfect competition among (potentially) differentiated traders. By matching the differences in treatment and control outcomes (such as prices and quantities), as well as the pass-through estimate, to their theoretical counterparts, we recover key market structure parameters.[1]

In the model, two parameters shape the competitiveness of the market: the degree of differentiation among traders and the market size (i.e., the number of competitors). First, we can recover the former from the difference in prices paid by treatment and control traders: in a market with perfectly homogenous traders, there should be one price, thus no difference between the two groups. In the experiment, treatment traders pay similar price (no significant difference) than control traders and they are more likely to provide advance payments to their regular suppliers (+14 percentage points, a 117% increase). The experimental results allow us to compute farmers rate of substitution between higher prices and a larger share of advance payments.

Once accounting for advance payments, the difference in the "effective price paid by treatment and control traders is about one-fifth of the subsidy value. The differentiation parameter we recover is then equal to 0.25 on a scale from 0 (i.e., perfectly homogenous

[1]As our theoretical framework in Section 3 discusses at length, the response of the control traders to increased competition from subsidized traders implies that the experimental differences between treatment and control outcomes cannot be interpreted as standard treatment effects. Nevertheless, for brevity purposes, we refer to these differences as "treatment effects in the rest of the paper.

traders) to 1 (i.e., each trader is a local monopsonist). The empirical result suggests that, in our context, differentiation comes from the fact that certain traders provide advance payments to farmers and they cannot be easily substituted in this service provision. This result is consistent with the significant literature on interlinked transactions in agriculture, which emphasizes the role of buyers in providing inputs and credit services.[2]

Second, we can infer the numbers of traders that compete for a given farmer. By treating it as an unknown parameter, we acknowledge that villages may not be the relevant definition of a market: farmers may sell, directly or indirectly, to traders that do not operate in their village or, on the other hand, they may sell only to traders that come to their farm gate. Given the level of differentiation estimated above, we find the number of traders consistent with the observed pass-through rate from wholesaler to trader prices (0.92). The "effective market size we recover is 15 traders —about twice the number of traders operating in a village (7.8). This finding speaks at the difficulty of defining market size ex-ante.

Overall, consistent with the survey by Dillon and Dambro (2016), the market is then quite competitive: in response to the subsidy, treatment traders —who are about one-fifth of the traders operating in the market—purchase substantially more cocoa than control farmers (+188%). The model suggests that the experiment increased (decreased) the quantity purchased by treatment (control) traders by about 118% (20%), relative to a counterfactual without the experiment, without basically raising the total quantity sold in the market. This experimental results is in line with the high pass-through rate from wholesaler to trader prices and with the low traders markdowns at baseline.

Finally, our paper also makes a methodological contribution. A number of recent papers have studied price effects of interventions such as cash transfers and farmer price support schemes using experimental designs that vary the *treatment saturation* at the cluster level typically a village or market (see, e.g., Baird et al., 2014; Cunha et al., 2015, Burke, 2014; Mobarak and Rosenzweig, 2014) The randomization saturation approach *assumes* the size of the cluster and that there are no cross-cluster externalities. In the approach of this paper, we instead move from assumptions over preferences and equilibrium concepts to *estimate* the market size (and other market structure parameters) using an individual-level randomization. In addition, we can still shed light on what would happen in experiments that treated a

[2]Bardhan (1980), Bardhan and Udry (1999) and Bell (1988) summarize a large body of theory looking at transactions that relate land, labor, product, and credit markets. Blouin et al. (2013), Casaburi and Willis (2015), Casaburi and Macchiavello (2016), Ghani and Reed (2014), and Macchiavello and Morjaria (2015a) provide primarily empirical contributions. This paper emphasizes the relationship between interlinked transactions and price transmission.

different share of traders in the market. We hope this approach will prove useful for the many cases where the cluster saturation approach is not feasible. This may be the case, for instance, if the size of the cluster (i.e., the market) cannot be easily determined ex-ante or if a plausible cluster would be too large to serve as unit of randomization.[3]

We emphasize that the methodology developed in this paper requires committing to a specific model of imperfect competition, which in our case is admittedly a very stylized one. We provide additional elements to mitigate concerns regarding this choice. First, we match other results of the experiment (i.e., *percent* differences between treatment and control in prices and quantities) to their theoretical counterparts. The estimates for the differentiation rate among traders and for the market size we obtain with this approach are very close to the previous ones. Second, we show that alternative models such as Bertrand and collusion among traders are not consistent with the data.

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 derives the equilibrium expressions and comparative statics expressions. Section 4 presents the results from the randomized controlled trial – including the estimation of the value of advance payments and the estimates of pass-through from the natural experiment. Section 5 combines the theoretical and empirical results to recover market structure parameters. Section 6 offers concluding remarks and discusses avenue for future research.

2 Study Setting and Experimental Design

2.1 The Sierra Leone Cocoa Value Chain

West Africa produces two-thirds of the world cocoa supply. Though given its small size Sierra Leone accounts for only a small share of this total, 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 industry has also grown tremendously in the last decade, with the value of exports growing ten-fold between 2009 and 2001, when the country's decade long civil war came to an end.

The within-country cocoa trade in Sierra Leone is highly fragmented across many traders, and the supply chain has many links, similar to other agricultural markets in developing

[3] A distinct line of experimental research studies the impact of providing price information to farmers (see, e.g., Fafchamps and Minten, 2012; Hildebrandt et al., 2015; Mitra et al., 2013).

economies (for examples in Africa see Fafchamps et al., 2005 and Osborne, 2005).[4] Farmers sell to traders, who sell to wholesalers in small towns, who in turn sell to exporters in larger towns, who in turn sell to buyers at the port. While the study of competition is surely relevant at each of these links in the supply chain, we focus on the final link closest to production, and leave the examination of other levels for future research.

The provision of loans by traders to farmers is an important characteristic of this industry. Traders will offer farmers credit before and during the harvesting season, which typically lasts from the beginning of the rainy season in July/August until the end of the year. Farmers use the advance payments for production (e.g. hiring workers for harvesting) or for consumption smoothing. Traders then allow farmers to repay the loan in cocoa by selling at a below market price for subsequent sales. The extension of credit may limit ex-ante and ex-post competition. First, only those traders who have enough information about a given farmer may offer credit. Second, once a farmer receives credit from a given trader, selling to other traders may be costly since customary authorities may step in to enforce contracts (see, e.g., Acemoglu et al., 2014 and Sandefur and Siddiqi, 2013) or because traders stop offering credit in subsequent seasons (see, e.g., Fafchamps, 2003 and Macchiavello and Morjaria, 2015b).

As emphasized by Atkin and Donaldson (2015), when looking at prices, it is important to focus on narrowly defined homogenous 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 measure prices, we worked with the partner wholesalers to refine a quality grading 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.[5] The analysis in this paper focuses on grade A

[4]Sierra Leone's cocoa industry is similar to those in Cameroon, Côte d'Ivoire and Nigeria all of which 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 war, 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.

[5]In Appendix A we discuss in greater detail the construction of the grades, and their relationship to wholesale prices and international standards of cocoa quality.

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.[6] These wholesalers collect cocoa in their warehouses, and then sell it on to exporters in the provincial capital of Kenema. Our sample includes 80 traders, henceforth study traders. This comprises almost all of the traders who do business regularly with these wholesalers.

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

Randomization occurred at the trader level. To improve the statistical power of our experiment, we implement a pairwise randomization strategy, first grouping traders in pairs and then allocating one trader to treatment in each of the pairs (see, e.g., Bruhn and McKenzie, 2009). We first match traders within wholesalers according to a self-reported estimate of the number of bags that they had sold since the beginning of the cocoa season, a plausible proxy for the scale of their business. Having matched the traders, we assigned treatment and control within pairs using a random number generator.[7]

2.3 Data Collection and Summary Statistics

Over the course of the experiment we collected a variety of data from traders.[8] At baseline, we interviewed each trader about his experience in the industry, and collected basic

[6]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 in 2011 joked with some cynicism that the cocoa stocks of the largest wholesalers in Pendembu could not come close to filling it.

[7]Of the 84 traders identified in the initial census, 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.

[8]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.

demographic indicators. We also asked traders to list each farmer they buy from regularly and all of the villages in which they buy. For each farmer, we asked whether the trader had given the farmer a loan over the past 12 months. Table 1, Panel A presents the results. Traders have sold an average of 19 bags of cocoa in the current harvest season. This suggests they operate at a small scale in terms of value.[9] Traders are experienced, with an average of 6.5 years selling to the wholesaler. Their average age is 38 years, 82% of the 46 year male life expectancy reported by World Health Organization in 2011. Traders are well off relative to the population. 58% have a cement or tile floor as opposed to dirt or thatch, a useful indicator of asset wealth in this context, and 92% own a mobile phone. The 2007 National Public Services survey of reports 18% and 8% respectively for ownership of these two assets among all households in rural areas. 83% of traders have access to a storage facility.[10] In the baseline listing, traders mention purchasing cocoa from 125 villages. The average trader operates in 4.6 villages, and buys from 25.9 farmers, an average 5.7 per village.[11] In the baseline survey, traders have given at least one loan to on average 70% of their clients. The third column of Table 1 shows that treatment and control groups are balanced on these trader-level covariates, according to regressions with randomization pair fixed effects.

During the experiment, when traders arrived at the warehouse, inspectors from our research team measured and documented the quality of their shipment. We collected 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 (we drop the entire weeks from the analysis). In these 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 from. 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 credit provision, in November and December we asked again the traders if they had given loans in the previous month to the farmers listed at baseline.

In the three weeks preceding the intervention, 56 of the 80 traders (27 control and 29 treatment) visited the warehouses. Table 1 Panel B shows that treatment and control groups

[9] A back of the envelope calculation suggests traders revenues since the beginning of the season was about $4,360.

[10] Field work suggests that traders for the most part do not perform long-term storage or other post-harvest processing activities.

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

are balanced along (weekly) quantity purchased and prices paid to farmers.[12]. 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 simple static framework of imperfect competition among buyers and models the impact of a subsidy to a subset of buyers (akin to our experimental treatment). We derive closed form solutions as a function of the primitive parameters for equilibrium prices and quantities and, crucially, for the differences in price and quantity between treatment and control buyers.

3.1 Setup

Producers

Consider measure one of homogenous producers operating in a market. Each producer produces output q. A producer profit from selling her output to n buyers is:

$$V(p_1^E, ..., p_n^E, q_1, ...q_n) = q_0 + \sum_{i=1}^{n} p_i^E q_i - C(q_1, ..., q_n) = q_0 + \sum_{i=1}^{n} p_i^E q_i - (\alpha \sum_{i=1}^{n} q_i + \frac{1}{2}\beta \sum_{i=1}^{n} q_i^2 + \gamma \sum_{j \neq i} q_i q_j), \tag{1}$$

where q_0 is the output that is not sold to traders (e.g., consumed, not harvested), q_i is the output sold to trader i. Each buyer i can pay producers before harvest (i.e. advance payment) or after harvest. Unit prices paid in the two periods are p_i^0 and p_i^1, respectively. We assume producers and buyers have a common, constant, rate of substitution between payments at different times, λ. Thus, we define the "effective price" paid by the buyer i, $p_i^E \equiv \lambda p_i^0 + p_i^1$.

The producer cost function, $C(\cdot)$, features "love for variety": for a total amount sold to buyers, costs are lower if sales are equally split across buyers. Specifically, we adapt the standard model of buyer quadratic utility and imperfectly competitive producers to our setting, which thus features producer quadratic cost function. $1 - \frac{\gamma}{\beta} \in [0, 1]$ is the measure

[12]The quantity regressions include all the traders in the sample, with quantity sets to zero if a trader does not visit the warehouses. Price regressions are only defined for traders who visit the warehouses. However, since all of the analysis is done within matched pairs, we would maintain internal validity on the subset of pairs where both traders visit the warehouses if extensive-margin attrition were the only source of selection. Section 4.1.1 discusses other potential selection concerns.

of differentiation among buyers. If this index is equal to zero, traders are homogenous. If it is equal to one, each buyer is a local monopsonist.

Assuming $q_0 > 0$, the inverse supply function for buyer i is then:[13]

$$p_i^E = \alpha + \beta q_i + \gamma \sum_{j \neq i} q_j \tag{2}$$

Buyers

There are n buyers in the market. Buyer i profit is

$$\pi_i = q_i(v_i - p_i^E), \tag{3}$$

where q_i is the quantity purchased, v_i is the ("effective") resale price, and p_i^E is the"effective price" defined above.[14]

We assume buyers are ex-ante homogeneous in the resale price v. The experiment introduces a subsidy, t, for a share μ of the buyers. Therefore, $v_i = v + t$ for treatment buyers and $v_i = v$ for control ones. Below we refer to variables for treatment (control) buyers with subscript T (C).

We assume *Cournot oligopsonistic competition*: each buyer sets quantities strategically, taking into account competitors' choices. In Section 5.2, we provide several arguments in support of this choice and discuss alternatives such as Bertrand competition and collusion.

3.2 Equilibrium

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

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

[14]v should be interpreted as the *net* price, after subtracting other buyer per-unit costs (which are assumed to be constant). Consistent with our study setting, we assume that traders do not receive credit from wholesalers and thus $v_i^E = v_i$.

Equilibrium Quantities and Prices

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

$$\begin{aligned} q_T &= \frac{\alpha(\gamma - 2\beta) + v(2\beta - \gamma) + t(2\beta + \gamma(-\mu n + n - 1))}{(2\beta - \gamma)(2\beta + \gamma(n-1))}; \\ q_C &= \frac{(2\beta - \gamma)(v - \alpha) - \gamma\mu n t}{(2\beta - \gamma)(2\beta + \gamma(n-1))} \end{aligned} \tag{4}$$

From the inverse supply functions in Equation 2, we then obtain equilibrium prices:

$$\begin{aligned} p_T &= \frac{\alpha\beta(2\beta - \gamma) + v(2\beta - \gamma)(\beta + \gamma(n-1)) + \beta\gamma\mu n t + t(\beta - \gamma)(2\beta + \gamma(n-1))}{(2\beta - \gamma)(2\beta + \gamma(n-1))}; \\ p_C &= \frac{(2\beta - \gamma)(\beta(\alpha + v) + \gamma v(n-1)) + \beta\gamma\mu n t}{(2\beta - \gamma)(2\beta + \gamma(n-1))} \end{aligned} \tag{5}$$

It can be shown that prices imply variable markdowns, $\frac{p_i^E}{v_i}$. Treatment (control) quantities are increasing (decreasing) in the treament amount t and both are decreasing in the share of treated buyers μ. Both control and treatment prices are increasing in both t and μ. These (intuitive) comparative statics suggest that treatment changes the behavior of control firms, as well as control ones. In terms of the experiment, this strategic response represents a violation of the Stable Unit Treatment Value Assumption (SUTVA).

"Treatment Effects"

This subsection obtains the "treatment effects" of the experiment for prices and quantities as a function of the parameters capturing differentiation among traders in the market, $\frac{\gamma}{\beta}$:

$$\Delta p^E \equiv p_T^E - p_C^E = \frac{t(1 - \frac{\gamma}{\beta})}{2 - \frac{\gamma}{\beta}} \tag{6}$$

and

$$\Delta q \equiv q_T - q_C = \frac{t}{2 - \frac{\gamma}{\beta}} \tag{7}$$

Δp^E is *decreasing* in γ/β: if goods are homogeneous, control traders must charge a price closer to the treated ones to stay in the market.[15] On the other hand, Δq is *increasing* in γ/β: if goods are homogeneous, the treatment traders can expand more by stealing from control

[15] If $\gamma = 0$, i.e., each buyer is a local monopsonist, the linear supply function implies that each monopsonist passes one-half of the subsidy, that is $\Delta p^E = t/2$.

traders. Both treatment effects are increasing in the value of the treatment, t, but do not depend on the share of treatment traders, μ. An extra treatment trader takes away quantity from both control traders and other treatment traders, thus the impact on the difference between the two (i.e. the "treatment effect") is ambiguous. Given the specific functional form we adopt, it is zero.

Pass-Through

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

$$\rho \equiv \frac{\partial p_C^E}{\partial v} = \frac{\partial p_T^E}{\partial v} = 1 - \frac{1}{2 + \frac{\gamma}{\beta}(n-1)}, \tag{8}$$

which is, intuitively, increasing in γ/β and n, as both capture the competitiveness of the market.

3.3 Discussion

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 relies on repeated interactions. Second, we shut down any source of heterogeneity, aside from the one introduced by the experiment. Third, we rely on specific functional forms. For instance, the quadratic cost function leads to a linear supply function that is particularly easy to work with analytically but that cannot be microfounded with a discrete choice problem (see, e.g., Jaffe and Weyl, 2010 and Armstrong and Vickers, 2015). Fourth, we assume the experiment does not impact the number of players in a market as given, i.e., we don't allow for entry or exit.

The main benefit of this approach is that it delivers a clear relation between the treatment effects and the parameters of the model. In Section 5.2, we provide some evidence in support of the specific choice of the model. We leave to future research to enrich the framework by relaxing some of the assumptions discussed above.

4 Results

This Section presents the empirical results of the paper. First, it reports the impact of the experimental treatment on prices traders pay to the farmers, advance payment provision, and quantities purchased. Second, it uses cross-sectional and experimental variation to quantify the value of advance payments. Finally, using exogenous changes in the world price of cocoa, it estimates the pass-through rate of industry-wide price shocks

4.1 Experimental Results

We report the comparison of outcomes between treatment and control traders. As discussed above, and as the theory in Section 3 makes explicit, the differences in outcomes cannot be interpreted as treatment effects in the standard Rubin sense: the subsidy affects the behavior of both treatment and control traders, as they compete for suppliers. Rather, these differences will later be matched to their theoretical counterparts to recover the competition parameters of the model.

4.1.1 Prices

We focus on prices first. Figure 2 shows weekly averages for the duration of the study: i) world price (right y axis); ii) wholesaler price; 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 display two key features. First, trader prices follow closely wholesaler prices, which themselves 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 treatment and control traders pay to the farmers.

To study more in the detail treatment effects in prices, we estimate the following regression, where an observation is a shipment s delivered by trader i of randomization pair p, in week t:

$$Price_{sit} = \eta_p + \eta_t + \pi^p(\text{Treat}_i) + \epsilon_{sit}, \tag{9}$$

where η_p and η_t are randomization pair and week fixed effects, respectively.[16] We cluster

[16]Since pairs were matched within wholesalers, this effectively controls for the town in which the trader sells his cocoa.

standard errors at the trader level (the unit of treatment).[17]

The term π^p is the coefficient of interest.[18] Table 2, Column (1) presents a regression without week fixed effects: the coefficient is $\hat{\pi}^p = -32.5$ (se=47.16). In Column (2), we introduce week fixed effect and the coefficient becomes -5.4 (s.e. = 14.9). While the two coefficients are not statistically distinguishable from each other or from zero, the pattern suggests that selection in *when* to sell matters. In particular, it appears that the experiment induced treatment traders to stay longer in the market, as many of them continue to make purchases in December when the season was winding down and prices were low because of the international price decline.[19]

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 for trader characteristics and for the village where the majority of cocoa in the shipment originated.[20] The former group includes 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-level 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 pretty stable when including, trader controls, village controls and, finally, both set of controls. This suggests the selection concerns described above cannot drive the results. Overall, the various specification provide persuasive evidence that prices did not differ between treatment and control traders.[21]

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

[18]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. This provides reassurance that our price results are not driven by measurement error in prices.

[19]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.

[20]In the intervention period, eighty of the 125 villages listed at baseline appear as "main village in at least one shipment, covering approximately 85% of the suppliers listed at baseline.

[21]In results not presented, we also tested for effects on the prices of B and C grade cocoa. Though we found no significant effect on the price of grade C cocoa, we did find a statistically significant effect on grade B prices (the point estimate is 37, which is still very far from the value of the subsidy). Field interviews suggest that this result is a result of Type I error on the part of traders, who observe quality imperfectly. The bonus has increased the expected price for grade A quality cocoa relative to grade B. If quality is imperfectly observable, treatment traders may now be more willing to pay the grade A price premium for cocoa that has some probability of being grade A.

4.1.2 Advance Payments

To investigate the effect of the bonus on traders advance payments, we estimate the following regression:

$$AdvancePayment_{fip} = \eta_p + \pi_1^a(\text{Treat}_i) + \nu_{fi} \tag{10}$$

An observation is a farmer mentioned as a regular supplier in the trader baseline. $AdvancePayment_{fip}$ is an indicator of whether farmer f was given advance payments by trader i of pair p during the course of the experiment. The term π^a is the treatment effect estimator, and ν_{fi} is an error term.

Table 3 presents estimates of π^a. 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 treatment effect 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 12 percentage points. Columns (2)-(4) show that results are similar when adding trader controls, village controls, and both set of controls.

4.1.3 Quantities

Finally, we investigate the impact of the treatment on the quantities purchased. Figure 3 shows the weekly total 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 suggests that the experiment introduced an increase in the total quantity purchased by study-traders, consistent with the idea that treatment traders obtained cocoa from non-study traders. Finally, in the second part of the experiment, following the break described in Section 2, there is a stark reduction in total quantities purchased, suggesting that the season was essentially over by the time the intervention was ended.

Table 4 presents regression results from the following regression models:

$$Quantity_{ipt} = \eta_t + \eta_p + \pi^q(\text{Treat}_i) + \zeta_{it} \tag{11}$$

where an observation captures the total purchases of cocoa trader i, in randomization pair

p, makes in week t. During the experiment, treatment traders on average purchase 537 pounds per week more than control traders, roughly a 190% difference.[22] The results are robust when including trader controls in Column 2. Overall, this is a large impact of the treatment.[23] 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 model developed in Section 3 and the analysis in Section 5 shed light on these patterns.

4.2 The Value of Advanced Payments

This Section measures the value buyers may provide to farmers through the provision of advance payments. To clarify the empirical steps, we must first introduce a slight modification in the model to reconcile the theory with the available data. In the model, the effective price paid by buyer i was defined as $p_i = \lambda p_i^0 + p_i^1$. However, in the data, we do not observe the amounts of payments made at different times. Rather, we just observe whether a certain farmer receives advance payments or not. Therefore, we modify the definition of the (average) effective price paid by buyer i as it follows

$$p_i^R = p_i + \tilde{\lambda} \cdot ShareAdvances_i, \tag{12}$$

where p_i is the total monetary amount paid by the trader, $ShareCredit_i$ is the share of farmers to whom trader j provides advance payments and $\tilde{\lambda} > 0$ is the extent to which farmers (and traders) value advance payments credit (again, this is assumed to be constant).

We make progress in measuring the value of $\tilde{\lambda}$, using two sources of variation: cross-sectional and experimental. Table 5 presents these results from the cross-sectional analysis. Each column shows the coefficients from regression of a baseline village-level outcome on other village level covariates and a constant.[24] These indicators come from the pre-treatment

[22] We include all the eighty traders in the sample, assigning value zero to trader-week pairs with zero purchases, including for the six traders than never showed up during the experimental period). Results are similar when dropping these traders.

[23] The treatment effect on quantities is significantly 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.

[24] 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 in this section 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.

period, when our inspectors were collecting data on quantity, prices, and quality of cocoa delivered to wholesalers, but treatments had not been assigned. Table 5 Column 1 shows that moving from zero credit share to full credit share decreases the price paid (conditional on quality) by 137 Leones (s.e. = 74.6).[25] In terms of our model, that a trader paying a certain price and another trader paying a 137 Leones lower price but providing advance payment are offering the same *effective* price to the farmer.

We next use variation created by our experiment. We modify Equation (10) to allow for heterogeneous treatment effects across villages by specifying the regression equation:

$$AdvancePayment_{fipv} = \eta_p + \pi^a(\text{Bonus}_i) + (\text{Bonus}_i \times \mathbf{W}'_v)\boldsymbol{\pi}^{\boldsymbol{a}}_{\boldsymbol{w}} + \mathbf{W}'_v\boldsymbol{\beta}_{\boldsymbol{w}} + (\text{Bonus}_i \times \mathbf{X}'_{\mathbf{i}})\boldsymbol{\pi}^{\boldsymbol{a}}_{\boldsymbol{x}} + \mathbf{X}'_{\mathbf{i}}\boldsymbol{\beta}_{\boldsymbol{x}} + \nu_{\mathbf{fipv}} \tag{13}$$

where, as before, $\mathbf{W}_v$ is the vector of village covariates and $\mathbf{X}_i$ is a vector of trader level covariates. For any trader-village pair iv then we have an estimator for the credit response $\widehat{TE}^c_{iv} = \mathbf{W}'_v\boldsymbol{\pi}^{\boldsymbol{a}}_{\boldsymbol{w}} + \mathbf{X}'_i\boldsymbol{\pi}^{\boldsymbol{c}}_{\boldsymbol{i}} + \pi^c$. Finally, we run the following specification to test whether pairs with higher credit pass-through display lower price treatment effect:

$$Price_{kipvt} = \eta_p + \eta_t + \pi^p(\text{Bonus}_i) + \pi^p_a(\widehat{TE}^a_{iv} \cdot \text{Bonus}_i) + \mathbf{X}'_i\boldsymbol{\beta}_{\boldsymbol{x}} + \mathbf{W}'_v\boldsymbol{\beta}_{\boldsymbol{w}} + \epsilon_{kiptv}, \tag{14}$$

Substitutability across pass-through margins predicts that $\pi^p_a < 0$.[26]

Figure 4 provides initial evidence. We estimate price and advance payments treatment effect in each of the five chiefdoms included in the study, and plot them against each other. The scatter suggests a negative relation: a regression line has a slope of -271. We note that this is larger (in absolute value) than the value we estimated in the cross sectional analysis in Table 5. One likely explanation is that the baseline advance payment variable captures whether the farmer had received advances in the twelve months before the baseline, a much longer time horizon that the one of the experiment (two months). Thus, the two advance payment dummies may capture different intensities of payments. An ideal data collection would have gathered transaction-level information on advance and harvest payments.

Table 6 presents estimates of θ^p_c. In the different columns we show estimates generated using different sets of controls to predict $\widehat{TE}^c_{iv}$. Any test of significance in equation (14)

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).

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

[26]We note that the estimated village-level treatment effect on credit, $\widehat{TE}^a_{iv}$, is collinear with the vector of controls and thus cannot be included in the estimating equation.

must account for prediction error in the treatment effect on credit. To do so, we follow the recommendations of Bertrand et al. (2004) and Cameron et al. (2008), and present p-values for a Wald test against the null hypothesis that $\theta_c^p = 0$ calculated using the pair cluster bootstrap-t procedure of Efron (1981) Bootstrap clusters are defined at the village level.[27]

Our estimates of π_a^p are negative and statistically significant at conventional levels in each of the specifications. In column 1, $\widehat{TE}_{iv}^{aN}$ is predicted using only chiefdom dummies. Chiefdoms are local geographic units of legal and political administration, and, as discussed in Acemoglu et al. (2014) a plausible proxy for variation in contract enforcement institutions. The estimate using these dummies predicts that a village where the bonus raised the likelihood of credit provision to farmers by 14 percentage points—the mean treatment effect in Table 3—would display a price response 46.2 Leones lower than a village with no effect of the bonus on credit. This is economically relevant as it accounts for a reduction of pass-through by about 1/3 relative to perfect pass-through.[28] We find similar results in column 2, where the effect on credit is predicted using chiefdom dummies and village covariates, and in column 3, where we add trader covariates into the prediction of the credit effect. While the precise magnitude of these results may vary across specifications, the core result is confirmed: price and credit pass-through are substitutes. In villages in which traders respond to a wholesale price increase by raising credit supply to the farmers, price pass-through will be lower.

Having obtained estimates for the value of advance payments —λ in the model —, it is then possible to estimate treatment effect on the effective price, $\Delta p^R = p_T^R - p_C^R = \hat{\pi}^p + \hat{\hat{\lambda}} \cdot \hat{\pi}^a$.

4.3 Pass-Through of Industry-Wide Price Shocks

In the final empirical contribution of the paper, we study how the prices traders pay to farmers respond to changes in the prices wholesalers pay to traders. For this purpose, we rely on variation in wholesaler prices and we seek to identify plausible sources of exogenous variation. As a starting point, Figure 2, discussed in Section 4.1.1, showed a stark reduction in prices (around 22%) between the first and the second half of the experimental period.

[27]Specifically, we first estimate $\hat{\theta}_c^p$ on the full sample and generate a T-statistic, T_0 from a Wald test of the null hypothesis that $\theta_c^p = 0$. We then draw with replacement a sample of villages 1,000 times. For each draw, we predict $\widehat{TE}_{iv}^{cN*}$ and then use it to estimate another $\hat{\theta}_c^{p*}$, where the star indicates the bootstrapped sample. We then generate a test statistic T^* from a Wald test of the hypothesis that $\hat{\theta}_c^{p*} = \hat{\theta}_c^p$ using the standard error from the bootstrap estimate. We hold this test statistic in memory. After 1,000 draws, a p-value is calculated from (twice the) position of T_0 in distribution of test statistics T^*.

[28]Recall that perfect pass-through of the bonus would imply pass-through of 150.

Table 7 presents the results. Column (1) reports a basic OLS regression of the trader price vs. the wholesaler price, with no constant included. Throughout the table, standard errors are clustered at the level of the day. The coefficient is estimated at 0.91. This suggests a high-level of pass-through from traders to farmers. We notice that results are not directly comparable with studies that focus on pass-through of *border* prices to final producers/consumers, such as Atkin and Donaldson (2015).[29]

One obvious concern is that the change may be correlated with local supply shocks. We implement two strategies to mitigate it. First, we instrument wholesaler prices with the international future 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 to assume 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=14,010). 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 significantly weaker when adding village fixed effects).[30]

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. After introducing the methodology, we present the results and discuss their implications.

5.1 Methodology

We aim at recovering several competition parameters of the model by matching the experimental and quasi-experimental results presented in Section 4 to their theoretical counterparts derived in Section 3. The key insight is that, through the model, the empirical results can

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

[30]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.

shed light of the degree of competition in the market even if, in experimental terms, the SUTVA fails.

As a preliminary step, we need to introduce two additional assumptions to reconcile the theory with the available, imperfect, data. The first one concerns the behavior of traders that are not included in the study. In the model, traders are homogenous aside for the treatment status. Randomization ensures this is a reasonable assumption (in terms of averages) for traders included in the experimental study. However, the market also features traders that were not included in the study. On average, traders report that there are 7.8 traders in a village but only 3.2 study traders. This suggests that about 60% of the traders are non-study traders. We do not have data on them. This is an important limitation of the study. We then must make the assumption than non-study traders behave as the control traders in our study.

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, these 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). This may not be a particularly realistic assumption because our experiment shows that advance payment provision does respond to changes in the resale price.footnoteAppendix **??** provides a model of general symmetric imperfect *producer* competition, in the spirit of Weyl and Fabinger (2013) and Atkin and Donaldson (2015), allowing buyers to pay producers at two different times (e.g. spot payments and trade credit) and it derives a formula for the pass-through rate of an industry-wide shock in this case. In light of this consideration, our estimate of ρ (0.92) may be considered a lower bound. Similarly, we must assume that the advance payments to *regular* suppliers are representative of advance payments to all the farmers. This suggests our estimate of the advance payment treatment effect (and thus of the effective price one) may be an upper bound.[31]

The analysis focuses on two key set of parameters. First, the differentiation ratio, $1-\beta/\gamma$. Second, the "market size, n. While taking the degree of differentiation as a parameter to recover is a standard choice, the fact that we also consider the market size as an unknown variable deserves further explanation. Villages would seem a natural starting point to define

[31]The discussion below makes it clear that underestimating the pass-through rate would lead to an under-estimation of the market size parameter, n, and that overestimating the effective price treatment effect would lead to an over-estimation of the differentiation rate, $1-\frac{\gamma}{\beta}$.

a market. Most villages have a meeting point where sales occur. However, farmers may also sell, indirectly or directly, outside the villages. For instance, some farmers may bring their cocoa to other locations directly or through other farmers. In addition, certain market places may serve multiple villages. On the other hand, farmers also sell at the farm gate. These arguments suggest that the definition of a market may be a complicated process. For this reason, we choose to take the market size as a parameter to recover, rather than assuming its value.

By matching the theoretical expressions in Equations 6 and 8 to the empirical results, we obtain a system of two equations in two unknowns, β/γ, and n:

$$\begin{cases} \Delta p^E &= \hat{\pi}_1^p + \hat{\tilde{\lambda}} \cdot \hat{\pi_1}^a \\ \rho &= \hat{\rho} \end{cases} \tag{15}$$

We use the values for $\hat{\pi}_1^p, \hat{\pi_1}^a, \hat{\tilde{\lambda}}$, and $\hat{\rho}$ that we estimated in Section 4. Specifically, we use the median of the estimates across the different specifications reported in the tables: $\hat{\pi}_1^p = -5.81, \hat{\pi_1}^a = .14, \hat{\tilde{\lambda}} = 256.49$, and $\hat{\rho} = .92$. Substituting the previous results, Equation 15 can then be rewritten as:

$$\begin{cases} \frac{t(1-\gamma/\beta)}{2-\gamma/\beta} &= 30.09 \\ 1 - \frac{1}{2+\frac{\gamma}{\beta}(n-1)} &= 0.92 \end{cases} \tag{16}$$

Additional moments could be obtained by matching the theory expression for the *levels* of price (Equation 5) to its empirical counterpart. We opt not to use this moment and to rely only on experimental or quasi-experimental estimates. For a given value of the subsidy t—Leones 150,—the first equation (the experimental treatment effect on prices) is sufficient to recover the differentiation parameter ratio $1 - \gamma/\beta$.[32] The pass-through equations then is needed to recover n. Importantly, the estimation does not use the quantity treatment effect in the specific market because that would require a definition of market size, which we take as an estimand.[33]

[32]The estimation of the differentiation degree and of the market size is invariant to currency unit choice.

[33]See discussion in Section 5.2 and in Appendix B about a different method that uses *percent* treatment effect differences in prices and quantities. The use of percent differences implies again no need for market size ex ante definition.

5.2 Results

Solving the above system delivers the following results: *i)* the estimated value of the differentiation rate $1 - \frac{\gamma}{\beta}$ is 0.251; *ii)* the market size parameter, n, is estimated to be 15.02. The response in advance payments plays an important role in generating this differentiation. Repeating the exercise by setting the treatment effect on advance payments to zero and using the *upper bound* of the treatment effect on total payments reduces the differentiation rate to 0.185. Farmers may not be able to readily substitute advance payments from one trader to another. The provision of advance payments may thus be a source of trader differentiation because it depends on the extent of previous relationships or on the specific information and monitoring ability of the trader.

The estimated value of n is 15.02. This is about twice the average number of traders (study and non-study) 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.[34]

From the estimates of $1 - \beta/\gamma$ and n, one can obtain the parameters of the direct supply function. In particular, the ratio $(n-1)c/b$ is equal to 0.976. This suggests that if buyer i raises his price by 1, the increase in the quantity he purchases comes primarily (97.6%) from reduction in the volumes purchased by other buyers, as opposed to an increase in the aggregate quantity supplied by farmers.[35] This is consistent with the idea that, as the experiment was implemented at harvest time, farmers had little way to increase their supply in response to the price changes. Assuming that the total quantity in the market did not change, the experiment increased (decreased) the quantity purchased by treatment (control) traders by about 108% (27%), relative to a counterfactual without the experiment.[36]

Overall, the results suggest fairly competitive markets. This is in line with the review by Dillon and Dambro (2016) It 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 wholesale

[34]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 trades in the village in price, advance payments, and quantities. However, this is partially due to low power.

[35]This suggest the total quantity in the market is roughly constant. However, since study traders have a higher share of traders relative to non-study traders, the total quantity for study trader increases, consistent with Figure 3

[36]Treatment traders are about one-fifth of the total number of traders and, during the experiment, purchases roughly three times as much as the control ones. Without the experiment (and assuming total quantity is constant), all buyers would buy the same amount, i.e., 138% of what control traders purchase with the experiment. The numbers in the text follow.

is on average 0.92. Therefore, 0.08 is an upper bound on the markdown buyers obtain in the transaction (i.e., the markdown may be lower because of transport and other transaction costs).

The approach derived in this Section *assumes* a specific model, which is very stylized (as discussed in Section 3.3). 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 of the data to recover the same parameters of interest.[37] We derive theoretical expressions for the *percent* difference in prices *and quantities* between treatment and control traders and match them to their empirical counterparts (0.01 and 1.88, respectively). Using again the pass-through rate ρ as an additional moment, we recover the differentiation rate, $1 - \gamma/\beta$, the market size, n, and the intercept parameter, α (this latter is only identified up to a monetary unit parameter). The estimates are very similar to the previous ones: $1 - \gamma/\beta = 0.258$ and $n = 15.14$. The fact that the results are so close is not mechanical since the first, more parsimonious, approach uses *level* differences while this second method uses percent differences both in prices and in quantities. Indeed, Appendix Figures B.1 and B.2 show that how the estimates of the two parameters would be very different when using other arbitrary values for the treatment-control price percent difference in a neighborhood of the real ones.

Second, we discuss alternative models. While our baseline model is Cournot, we also implemented the steps describe above using Bertrand competition (while retaining other assumptions on producers and buyers). The procedure delivers unrealistic values (a negative γ/β 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. Quantity constraints (arising, for instance, from transport technologies) may be relevant in this setting. Another candidate model could be collusion: the small price treatment effects and high quantity effects may be consistent with treatment and control buyers forming a cartel. However, collusion is inconsistent with the high pass-through rate to industry-wide price shocks presented in Section 4.3 and with the low trader markdowns we discussed above.

As a final exercise, having gained confidence in the choice of the model, we use our results to inform what would happen if one replicated the experiment changing the number of traders that are treated (with a per-unit subsidy of one Leone). Figure 5 shows the increase in the

[37] Appendix B includes the derivations.

treatment and control prices, relative to the scenario without any experiment, as a function of the share of treated traders, $\mu \in (0, 1)$. In the specific experiment we ran, study traders were about 40% of all the traders and thus the treatment traders (half of the study-traders) were one-fifth of all the traders (i.e. $\mu = 0.2$). Our results suggest that, given this treatment share, treatment (control) traders should raise prices by 0.34 (0.14) per unit of subsidy to the treatment trader, relative to a benchmark market without the experiment. In summary, when interpreted through the lens of the model, results from individual-level randomizations can still shed light on what would happen if one ran experiments with a different number of traders in the same market. This is an important result to inform industrial policies targeting specific agents in the market. For instance, policy makers may be interested in studying the impact *on the famrers* of providing subsidies to parts of the trading industry.

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. In this paper, we have developed an experimental approach where we combine treatment effects of a subsidy intervention and a model of imperfect competition. By matching the empirical results to their theoretical counterparts, we recover key market structure parameters.

We obtain two key results. First, traders are differentiated. The differentiation parameter we recover is then equal to 0.25 on a scale from 0 (i.e., perfectly homogenous traders) to 1 (i.e., each trader is a local monopsonist). The empirical result suggests that, in our context, differentiation comes from the fact that certain traders provide advance payments to farmers and they cannot be easily substituted in this service provision.

Second, we can infer the numbers of traders that compete for a given farmer. By treating it as an unknown parameter, we acknowledge that villages may not be the relevant definition of a market: farmers may sell, directly or indirectly, to traders that do not operate in their village or, on the other hand, they may sell only to traders that come to their farm gate. Given the level of differentiation estimated above, we find the number of traders consistent with the observed pass-through rate from wholesaler to trader prices (0.92). The "effective market size we recover is 15 traders —about twice the number of traders operating in a village (7.8). This finding speaks at the difficulty of defining market size ex-ante.

Overall, consistent with the survey by Dillon and Dambro (2016), the market is then quite

competitive. The model suggests that the experiment increased (decreased) the quantity purchased by treatment (control) traders by about 118% (20%), relative to a counterfactual without the experiment, without basically raising the total quantity sold in the market. One important caveat is that lower levels of the value chains (e.g. wholesalers, exporters), may be substantially less competitive than traders.

Methodologically, our analysis suggests that it is possible to learn about markets from individual level randomization (as opposed to randomization occurs by varying treatment intensity at the market level). We hope this approach will prove useful for the many cases where the cluster saturation approach is not feasible. This may be the case, for instance, if the size of the cluster (i.e., the market) cannot be easily determined ex-ante or if a plausible cluster would be too large to serve as unit of randomization.

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.

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.

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.

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, 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*.

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."

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*.

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. 2015. "Time vs. State in Insurance: Experimental Evidence from Kenya Contract Farming."

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.

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.

Efron, Bradley. 1981. "Nonparametric standard errors and confidence intervals." *canadian Journal of Statistics*, 9(2): 139–158.

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.

Fafchamps, Marcel, and Bart Minten. 2012. "Impact of SMS-based agricultural information on Indian farmers." *The World Bank Economic Review*, 26(3): 383–414.

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.

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.

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.

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.

Osborne, Theresa. 2005. "Imperfect competition in agricultural markets: evidence from Ethiopia." *Journal of Development Economics*, 76(2): 405–428.

Sandefur, Justin, and Bilal Siddiqi. 2013. "Delivering justice to the poor: theory and experimental evidence from liberia." 20.

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

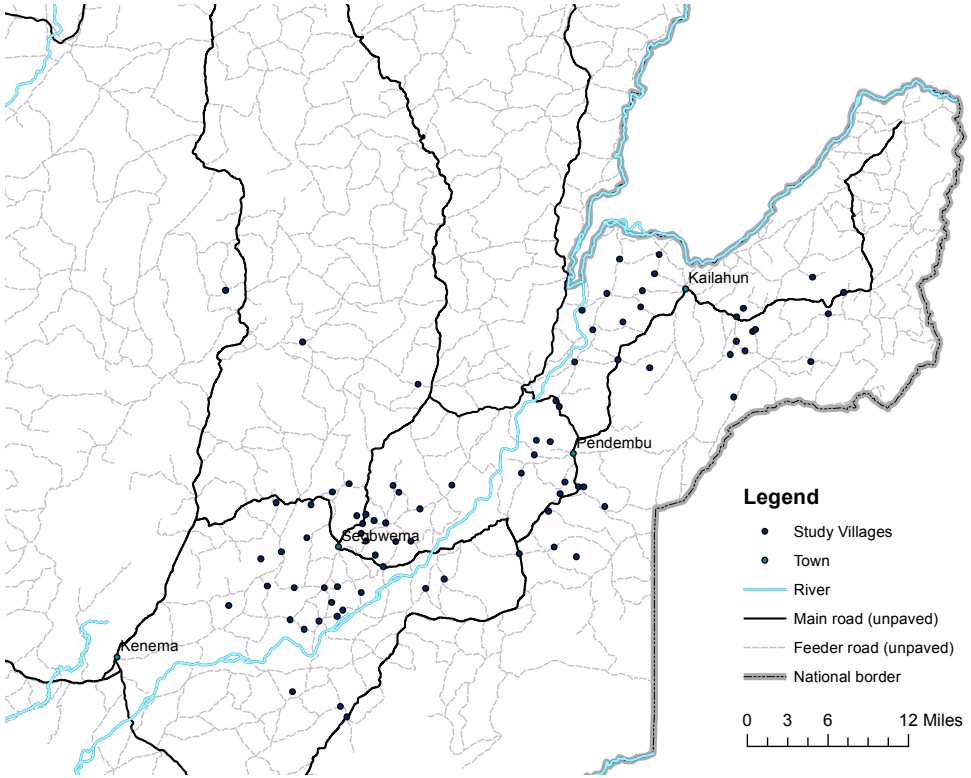
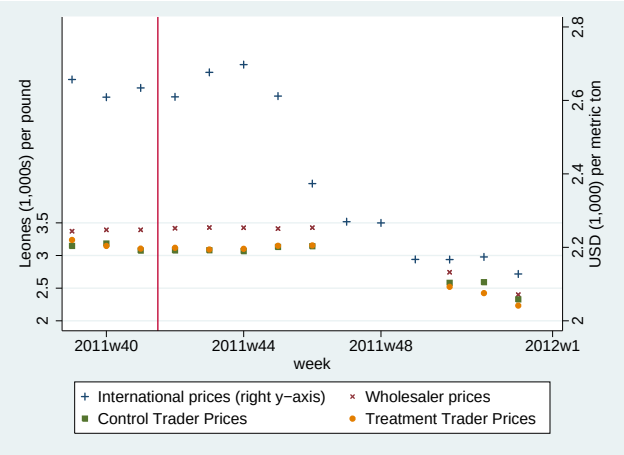
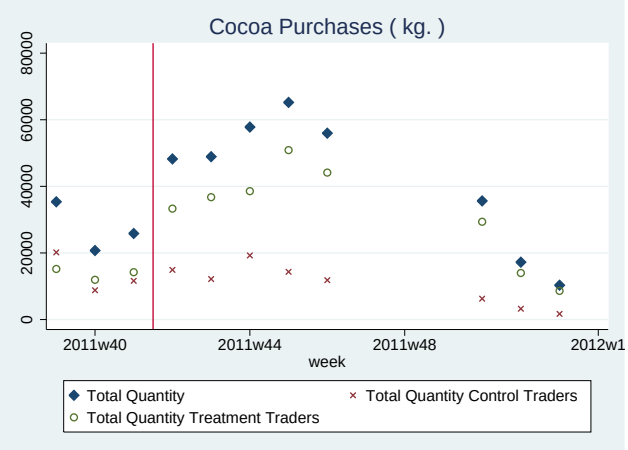


Figure 2: Cocoa Prices



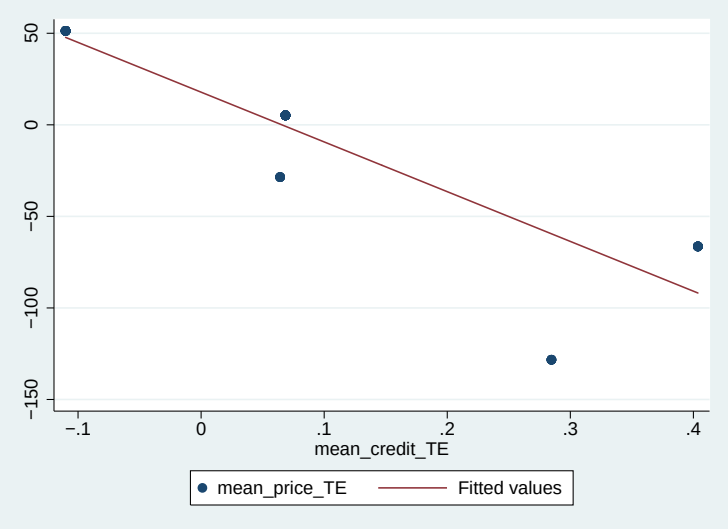
*N*otes: The figure presents average weekly prices for: *i)* international cocoa prices; *ii)* prices the study wholesalers pay to the traders; *iii)* prices control traders pay to farmers; *iv)* prices treatment traders pay to farmers. Wholesaler and trader prices data collection was suspended for most of three weeks (w47-w49). The vertical line marks the beginning of the intervention period.

Figure 3: Purchases of Cocoa



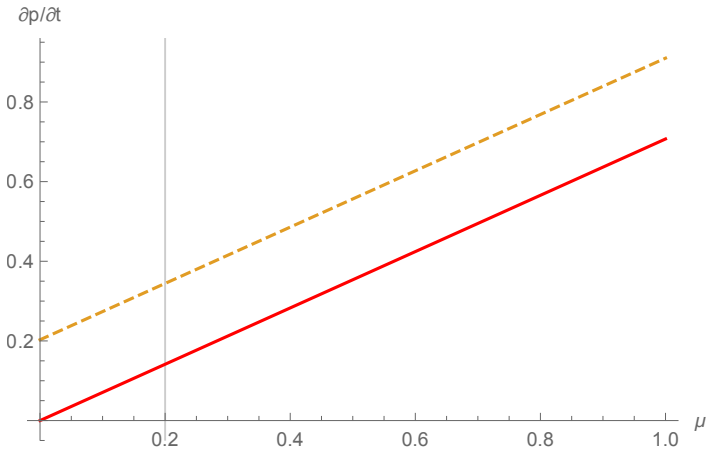
*N*otes: The figure shows the total amount of cocoa purchases by study (i.e., control+treatment) traders, control traders, and treatment traders. The vertical line marks the beginning of the intervention period. Data collection was suspended for most of three weeks (w47-w49).

Figure 4: Treatment Effects for Prices vs. Advance Payments by Chiefdom



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

Figure 5: Simulation: Treating Different Shares of Traders



*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 a function of the share of treated traders (out of the total number of traders in the market), μ. The vertical line reports the (approximate) share of traders treated in our experiment, $\mu = 0.2$. The graph shows that, at this level of treatment, control (treatment) traders would increase their effective prices by 0.14 (0.34), respectively, relatively to a scenario without the experiment. For $\mu \to 1$, the response of treatment traders tends to the pass-through rate, 0.91.

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	1.4
			(5.4)
Age, years	38.2	36.9	1.3
			(2.1)
Years selling to study wholesaler	5.7	7.3	-1.4
			(1.1)
Cement or tile floor in house $\in \{0, 1\}$	0.53	0.62	-0.09
			(0.1)
Mobile phone owner $\in \{0, 1\}$	0.90	0.93	-0.03
			(0.06)
Access to storage facility $\in \{0, 1\}$	0.88	0.78	0.10
			(0.09)
Villages operating in	4.87	4.25	0.62
			(0.41)
Number of farmers buying from	18.7	20.5	-1.8
			(3.1)
Share of farmers given credit since March	0.688	0.694	-0.006
			(0.0704)
Panel B: Pre-treatment shipment data			
Price Paid to Farmer (shipment-level)	3,086	3,114	8
			(20)
Pounds sold during pre-treatmen (weekly)	338	344	6
			(62)

Notes: The differences between treatment and control come from regressions with randomization pair fixef effects. Standard errors are clustered by trader.

Table 2: Price Responses to the Experiment

	(1)	(2)	(3)	(4)	(5)
Treatment Trader	-32.52	-5.47	-5.96	-11.69	-5.81
	(47.16)	(14.95)	(16.94)	(12.99)	(15.28)
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	1079	1079	1079

Notes: The table reports the difference between the prices paid by treatment and control traders to farmers during the experiment, measured in Leones per pound. The subsidy to treatment traders was Le. 150. per pound. An observation is a shipment delivered by the trader to a wholesaler. Trader controls are baseline values of pounds of cocoa sold, number of villages operating in, number of suppliers buying from, share of clients given credit in baseline, age, years of working with wholesaler, and dummies for ownership of a cement or tile floor, mobile phone and access to a storage facility. Village controls are baseline share of suppliers begin given credit, number of other bonus traders and number of study traders, miles to nearest town, and number of clients across all traders. 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***	0.14***	0.13***
	(0.03)	(0.02)	(0.03)	(0.02)
Control Group Mean	0.12	0.12	0.12	0.12
Trader Controls		X		X
Village Controls			X	X
Observations	1915	1915	1915	1915

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. 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	640

Notes: The table reports the difference between the quantities of cocoa purchases 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. Standard errors are clustered at the level of the trader. ***p<0.01, **p<0.05, *p<0.1.

Table 5: Substitutability, baseline correlations

Variables	(1) Baseline credit share	(2) Price per lb.	(3) Share of grade A
Pounds of cocoa per farmer	0.09***		
	(0.03)		
Number of study traders	0.03**	-3.8	-0.02*
	(0.01)	(7.6)	(0.01)
Miles to nearest town	0.00	-3.3	0.00
	(0.00)	(2.9)	(0.01)
Baseline credit share		-137.0*	0.34**
		(74.6)	(0.14)
R^2	0.07	0.09	0.11
Number of observations	125	44	75
Dependent variable mean	0.27	3,147	0.29

Notes: An observation is a village. Robust standard errors in parenthesis. The number of observations changes across columns because during the pre-treatment period, price and grade data are available for only the subset of the villages from which cocoa was delivered in that period. In particular, traders brought cocoa bags from only 75 villages and grade A cocoa from only 44 villages during the pre-treatment period. *** p<0.01, ** p<0.05, *p<0.1

Table 6: Substitutability of Pass-Through Margin

Variables	(1) Price	(2) Price	(3) Price
(Bonus × estimated effect of bonus on credit)	-330.09 [.014]	-256.49 [.006]	-182.71 [.008]
R^2	0.91	0.91	0.91
Number of observations	1,069	1,069	1,069
Chiefdom fixed effects	YES	YES	YES
Village controls	NO	YES	YES
Trader controls	NO	NO	YES

Notes: Each column presents estimates of θ_c^p from equation 14. P-values in brackets are derived from pairs cluster bootstrap-t at the village-level using 1,000 replications. Trader controls are baseline values of pounds of grade A 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. Trader and village controls are summarized in Tables 1 and **??**, respectively.

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 (during the three weeks before the intervention or the eight weeks of the intervention). In Columns (2)-(5), wholesaler prices are instrumented with the front-month prices for liquid cocoa futures, obtained by the *Intercontinental Exchange* (measured in US Dollars per 10 metric tons). Standard errors are clustered by day. ***p<0.01, **p<0.05, *p<0.1.

Appendix

A 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 see CAOBISCO (2002) and for a manual specific to West Africa on how to improve cocoa at the farm level see 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.

Table A.1 shows the average quality and wholesale prices of cocoa bags from the experiment, before the November fall in the international price. As can be seen, moisture content has the highest price elasticity—price falls by 0.32% with a one percentage point increase in moisture. Moisture is an important variable in our market, because wet cocoa rots in storage, destroying value. At an average 11% moisture content, cocoa in our market is substantially wetter than export grade, which requires a maximum moisture content of 7%. For this reason, many exporters maintain large drying facilities. There is an efficiency cost to this organizational structure, as some cocoa that is not dried at the farm gate will be lost to rot in transport.

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. At baseline, quantities supplied by traders were approximately one third of each.

Table A.1: Appendix, Cocoa Quality

Defect	Average per shipment	Price elasticity	Average per pound price by tercile of defect (Le.) 1	2	3
Moisture Content	11%	-0.32%	3,384	3,297	3,263
Mold	2%	-0.02%	3,308	3,353	3,241
Germinated	3%	-0.01%	3,309	3,313	3,298
Under-fermented	15%	-0.02%	3,345	3,333	3,228
Slate	7%	-0.01%	3,323	3,304	3,279

Notes: Data from 916 treatment and control transactions. Elasiticty gives the percentage reduction in price for a 1 percentage point increase in the defect.

B Recovering Competition Parameters: An Alternative Approach

This Appendix presents details about the alternative approach we use to recover the competition parameters $1-\gamma/\beta$ and n, which we briefly discussed in Section 5.2. Our goal is to identify different (quasi-)experimental moments and to compare the results obtained when using these moments to the ones of main approach presented in Section 5.1. Showing that different moments deliver similar estimates would provide support for the specific theoretical model we use.

B.1 Methodology

The main approach described in Section 5.1 relies on two moments: the *level* difference in treatment and control prices (Equation 6) and the pass-through rate of changes in wholesaler prices (Equation 8). In this section, we show how the key parameters $1-\gamma/\beta$ 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:

$$\%\Delta p^E \equiv \frac{p_T^E - p_C^E}{p_C^E} = \frac{t(1-\frac{\gamma}{\beta})(1+\frac{\gamma}{\beta}(n-1))}{\frac{\gamma}{\beta}\mu n t + (2-\frac{\gamma}{\beta})(\frac{\gamma}{\beta}(n-1)v + (\alpha+v))} \tag{B.1}$$

and

$$\%\Delta q^E \equiv \frac{q_T^E - q_C^E}{q_C^E} = \frac{t(-2-\frac{\gamma}{\beta}(n-1))}{\frac{\gamma}{\beta}\mu n t - (2-\frac{\gamma}{\beta})(v-\alpha)} \tag{B.2}$$

For a given value of the subsidy t, these expressions depend on additional parameters, i.e., μ, v, α, as well as on those we aim to recover, i.e., γ/β 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, such as transport and storage costs. We set $v = 3,145$, which implies a 5% markdown.[38]

[38]Results are fairly stable when using alternative values of v between 2,987 and 3,260.

B.2 Results

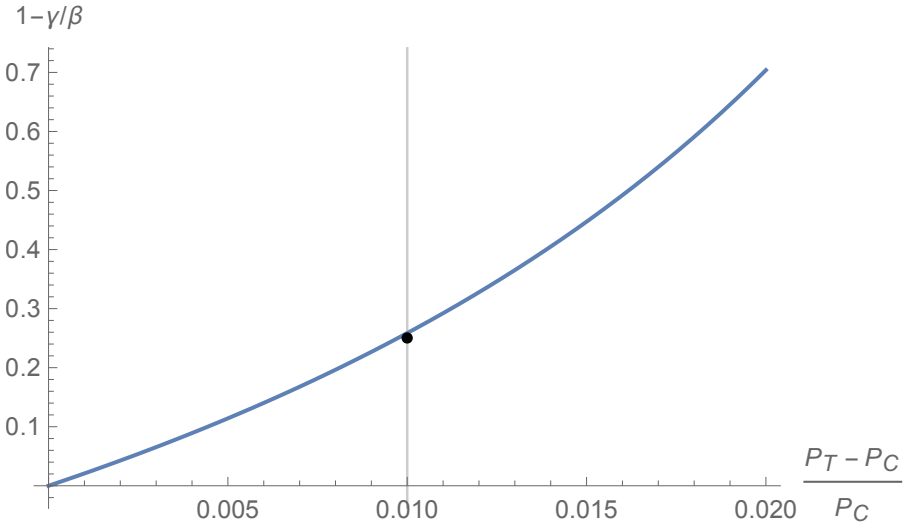
Having assigned values to μ and v, we have a system of three equations—Equations B.1 and B.2 defined above and the pass-through formula (Equation 8)—, in three unknowns, $\frac{\gamma}{\beta}, n$, and α. We note that the intercept term α is identified only up to the currency unit choice.-

When solving the system, the value of the three parameters is $\frac{\gamma}{\beta} = 0.742$ (and thus $1 - \frac{\gamma}{\beta} = 0.258$), $n = 15.14$, and $\alpha = 2085.75$. The results for $1 - \frac{\gamma}{\beta}$ and n are very close to the ones obtained when using the more parsimonious methodology described in the main text. We see this as evidence in support of the specific competition model chosen for the analysis.

Finally, we emphasize that the similarity of the results between the two approaches is not a mechanical result. This can be easily seen: one uses the *level* of the difference between treatment and control prices, while the other uses the *percent* the difference between treatment and control in both prices and quantities. Figure B.1 and B.2 confirm this point: the two graphs show, respectively, how the estimated values of $1 - \gamma/\beta$ and n would vary with different values of the percent treatment effect in prices, $\frac{p_T - p_C}{p_C}$, in a neighborhood of the real value, 0.01 (represented by the vertical gray line). In each graph, the large dot reports the estimate from the main approach describe in Section 5.1. The key take-away point is that, while the estimates 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 level price treatment effect were equivalent to a different percent price treatment effect).

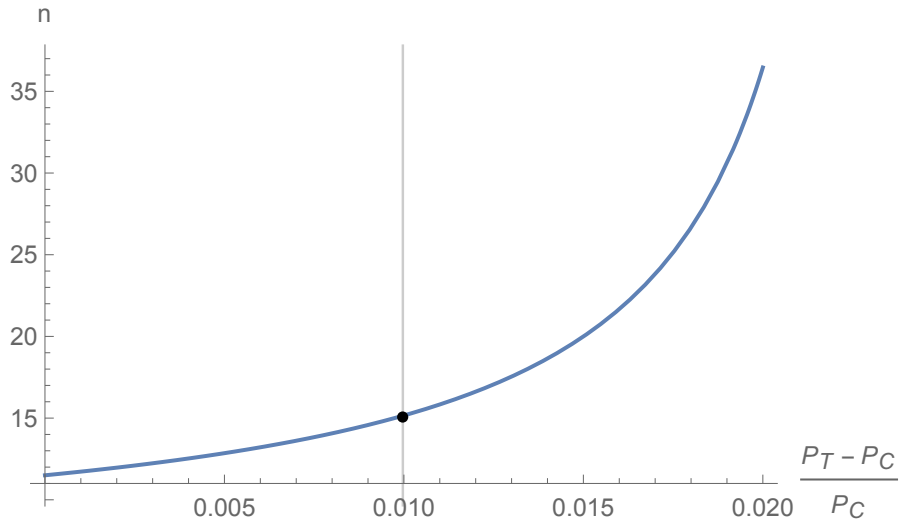
Appendix B: Figures

Figure B.1: Sensitivity of $1-\gamma/\beta$ to $\frac{p_T-p_C}{p_C}$



Notes:

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



Notes: