

Competition in Agricultural Markets: An Experimental Approach*

Lorenzo Casaburi Tristan Reed

February 2017

Abstract

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

JEL Classification: O13, O17, H2, F14.

Keywords: Agricultural markets, competition, pass-through, interlinked transactions, field experiment methodology.

*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, Robert Gibbons, Rachel Glennerster, Gita Gopinath, Oliver Hart, Asim Khwaja, Michael Kremer, Ted Miguel, Esteban Rossi-Hansberg, Orié Shelef, Tavneet Suri, Chris Udry, Jack Willis, and workshop participants at CSAE Oxford, Harvard, LSE/UCL, MIT, the Montreal Workshop on Productivity, Entrepreneurship and Development, NBER Development Meeting, NBER Development and Organizational Economics Workshop, Paris School of Economics, Stanford, Stockholm University, University of Naples, Trinity College Dublin, UC Berkeley, UC San Diego, 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 the study of agricultural markets in Sub-Saharan Africa and other low- and middle-income regions. Competition shapes how price signals propagate along supply chains, and thus the welfare implications of taxes and subsidies for producers and buyers. In a recent survey of the empirical literature on competition in Sub-Saharan Africa, Dillon and Dambro (2016) point out that the evidence is “remarkably thin” and that, where conclusions can be drawn, they point to a high degree of competition. 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 (Fafchamps et al., 2005 and Osborne, 2005), price dispersion (Fackler and Goodwin, 2001; Aker, 2010), or the pass-through of international prices along the value chain (Fafchamps and Hill, 2008). In this paper, we propose an experimental approach to estimate the degree of competition. Our design is based on the randomization of unit subsidies to traders for their purchases from farmers.

Randomization ensures orthogonality between the cost shock (i.e., the subsidy) and trader characteristics. Therefore, by comparing prices of subsidized and unsubsidized traders *in the same market*, we can recover the trader differentiation rate, which is the key market structure parameter in a standard model of imperfect competition. In intuitive terms, this parameter captures the extent to which competitor prices affect the supply of a trader, given her price.¹ The estimation uses only price data, thus addressing potential concerns the above methods may face in terms of costly data requirements and measurement errors. In addition, we combine our experimental results on prices and quantities with quasi-experimental estimates of the pass-through rate to estimate the market size (i.e., the number of traders competing in a market).

An ideal experiment aimed at measuring 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 the pass-through rate (i.e., the difference in prices between two markets with different producer costs) is a key object to recover the degree of market power. However, researchers face two common

¹More precisely, since we work with quantity-setting traders, the differentiation rate depends on the ratio of the slopes of the inverse supply to own vs. competitor quantities.

constraints in the design of market-level randomized trials. First, 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 to study market structure.

We start by presenting a standard framework of oligopsonistic competition among traders (see, e.g., Vives, 2001). In the framework, farmers sell their output to traders (at the “trader price”) and traders resell it to wholesalers (at the “wholesaler price”). We model the experiment by introducing a unit subsidy for a subset of traders on their purchases. Two parameters shape competition: the degree of differentiation among traders and the market size, defined as the number of traders farmers can sell to. We characterize equilibrium trader prices and quantities for subsidized and non-subsidized traders. We then derive the difference in these equilibrium outcomes between the two groups of traders. Finally, we derive the pass-through rate—i.e., the change in trader prices that occurs in response to a change in wholesaler price, common across all traders.

The model highlights the strategic interaction among subsidized and unsubsidized traders competing in the same market: unsubsidized traders adjust their choices in response to the subsidy of their competitors. Therefore, differences between treatment and control outcomes should not be interpreted as Rubin treatment effects. However, since randomization still ensures that subsidies are uncorrelated with firm characteristics, such differences nevertheless represent key moments to recover market structure parameters. Specifically, the difference between prices paid by (randomly) subsidized and unsubsidized traders depends on the differentiation rate: in a market with perfectly homogenous 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 in the market.² Similarly, the extent to which a subsidized trader can steal suppliers from unsubsidized competitors depends on differentiation, too.

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

We use the above insight to recover the differentiation parameter. In the experiment, treatment traders—about one-fifth of the traders operating in the study region—receive a per-unit subsidy worth about 5% of the average trader price. During the experimental period, they pay similar prices than control traders. They are also more likely to provide advance payments (i.e. liquidity services) to their regular suppliers (+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.³ Using two different strategies, we then compute the difference between treatment and control traders in the “effective price”—a price that accounts for the value of advance payments, akin to the net present value of the transaction. The difference in effective price 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.⁴

The experimental results are thus sufficient to recover the differentiation rate. This rate summarizes the level of competition in the model *for a given number of traders*. However, while market-level randomization assumes the number of traders competing in a market, we instead treat this variable as an unknown parameter. In doing so, we acknowledge that villages may not be the appropriate definition of a market: the number of traders observed to operate in a village may not correspond to the number of traders a producer could potentially sell to. Farmers may sell 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.⁵ We can estimate the number of traders by combing the above experimental results with quasi-experimental estimates of the pass-through rate from wholesaler to trader prices. 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—about 75% larger

³This observation is consistent with the recurring finding of the significant literature on interlinked transactions in agriculture, which emphasizes the role of buyers in providing inputs and credit services Bardhan (1980), Bardhan and Udry (1999) and Bell (1988) summarize a large body of theory looking at interlinked 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), Maitra et al. (2014), and Macchiavello and Morjaria (2015a) provide primarily empirical contributions. Our paper emphasizes the relationship between interlinked transactions and price transmission.

⁴This implies that the slope of a trader inverse demand to a competitor quantity is eighty to ninety percent of the slope to her own quantity.

⁵Sales 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.

than the number of traders observed operating in a village (7.8). This finding speaks at the difficulty of defining market size ex-ante. It also suggests that, despite high transportation costs, villages are still in part open economies.

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. In response to the subsidy, treatment traders purchase substantially more cocoa than control farmers (+188%). Through the model, we then quantify the impact of the experiment relative to a counterfactual scenario without the experiment. The experimental subsidy raises effective prices by about one-third of the subsidy (i.e. by about 1.5% of the average price level) for treatment traders and by one-sixth for control traders. Further, we find the difference in quantities purchased by treatment and control traders arise primarily from market stealing, as opposed to increases in aggregate supply, which we can bound at 0.8%. Specifically, the experiment increased the quantity purchased by treatment traders by about 108% and decreased that of control traders by 26%. The lack of change in aggregate quantity is consistent with the fact that the experiment was implemented halfway through the harvest seasons. By combining the price and quantity responses for control and treatment traders, we then compute that farmer revenues increased by 1.1 to 1.8%.

Finally, we consider the welfare implications for producers of counterfactual subsidy programs treating different shares of traders. In a country such as Sierra Leone, where population density is low and mobile money does not exist, subsidies paid to traders may be a cost-effective way of distributing income support to producers who could not otherwise be reached by government or aid agencies. We define the return on investment of these subsidies as the ratio between the increases in farmer revenues in response to the subsidy and the subsidy cost. In our experiment, which treated around 20% of the traders in the market, this return was about two-thirds of an intervention providing the same unit subsidy to *all* the farmers in the market. We also show that, for a given level of pass-through, the welfare implications of subsidies targeting only a subset of traders vary with the differentiation rate and the market size. This emphasizes the importance of estimating both competition parameters and not just their combination, which the pass-through rate summarizes.

Importantly, 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 use different moments (i.e., the *percent* differences between treatment and control in prices and quantities) to recover the two market structure parameters. The parameter estimates we obtain are very close to the previous ones (we also provide evidence that this result is not mechanical). In the spirit of an overidentification test, we interpret the fact that different moments generate similar estimates as evidence in support of the model. Second, we present evidence suggesting that alternative models such as Bertrand, monopsonistic competition à la Dixit-Stiglitz, 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 one of the first papers to use experimental subsidies to study 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. Rotemberg (2014) identifies the spillover effects on larger firms of a subsidy available only to small firms in India using variation in ineligible firms exposure to eligible competitors within a given administrative district. Busso and Galiani (2014) provides experimental evidence on the effect of increased competition on the prices and quality of goods, using randomized entry of new firms in retails markets in the Dominican Republic. Finally, 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.⁶ Relative to this body of work, our paper provides a proof of concept of how individual-level randomization can shed light on market-level features. We hope this approach will prove useful for the many cases where the market-level variation approach 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 derives the equilibrium and comparative statics expressions. Section 4 presents the results from the randomized controlled trial and the estimates of pass-through from the natural experiment. Section 5 combines the theoretical and empirical results to recover market structure parameters and to quantify the impact of the experiment. Section 6 offers concluding remarks and discusses avenues for future research.

⁶Other examples include Cunha 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).

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 the beginning of the rainy season in July/August until the end of the year.

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).⁷ Farmers sell to traders, who sell to wholesalers in small towns, who in turn sell to exporters in larger towns, who in turn sell to buyers at the port. While it is important to study the degree of competition in each of the links of the supply chain, we focus on the final 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 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 ensure we measure prices for a well-defined homogenous 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.⁸ The analysis in this paper focuses on grade A cocoa, the grade targeted by the

⁷Sierra 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 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.

⁸In Appendix A we discuss in greater detail the construction of the grades, and their relationship to wholesale prices and international standards of cocoa quality.

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.⁹ 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 until 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 matched 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.¹⁰

2.3 Data Collection and Summary Statistics

Over the course of the experiment we collected a variety of data from traders.¹¹ At baseline, we interviewed each trader about his experience in the industry, and collected basic

⁹These 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.

¹⁰Of 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.

¹¹One 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. 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 mention purchasing cocoa from 125 villages. The average trader operates in 4.6 villages, and buys from 25.9 farmers.¹² 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 5.2, 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.¹³ Contracts are enforced through customary authorities (see, e.g., Acemoglu et al., 2014 and Sandefur and Siddiqi, 2013) or through relational contracts (see, e.g., Fafchamps, 2003 and Macchiavello and Morjaria, 2015b). Study traders report having given at least one loan to about 70% of the suppliers listed at baseline.

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 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 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

¹²Figure 1 presents a map of the study setting.

¹³Interviews 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).

are balanced along (weekly) quantity purchased and prices paid to farmers.¹⁴ 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 for treatment and control traders and, crucially, for the differences in outcomes between the two groups.

3.1 Setup

Producers

The economy is composed by V villages. In each village, there is measure one of homogeneous producers, each producing output q , and n buyers compete for these producers. The inverse supply buyer i faces in a village is:¹⁵

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

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.

Following the literature (see, e.g., Vives, 2001), we define the *differentiation rate* $\Gamma \equiv 1 - \frac{\gamma}{\beta}$. If $\Gamma = 0$, buyers are homogeneous: the slope of the inverse demand to own quantity equals the slope to a competitor's quantity. If $\Gamma = 1$, buyers are local monopsonists: a buyer's inverse demand does not depend on other buyers' quantities. The differentiation rate Γ is

¹⁴The quantity regressions include all the traders in the sample, with quantity set to zero if a trader does not visit the warehouses. Price regressions include only traders who visit the warehouses. However, since all of the regressions include randomization pair fixed effects, 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.

¹⁵This inverse supply can be derived assuming the representative producer cost function features *love for variety*. Specifically, the inverse demand can be derived from the following producer profit function: $V(p_1^E, \dots, p_n^E, q_1, \dots, q_n) = q_0 + \sum_{i=1}^n p_i q_i - C(q_1, \dots, q_n) = q_0 + \sum_{i=1}^n p_i q_i - (\alpha \sum_{i=1}^n q_i + \frac{1}{2} \beta \sum_{i=1}^n q_i^2 + \gamma \sum_{j \neq i} q_i q_j)$, where q_0 is the output that is not sold to traders (e.g., consumed, not harvested) and q_i is the output sold to trader i .

one of the key market structure parameters we later aim to recover.¹⁶

Buyers

Buyer i profit in a given village, π_i , 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.¹⁷

We assume buyers are ex-ante symmetric in the resale price v .¹⁸ 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 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 5.3, we provide several arguments in support of this choice and discuss alternatives such as Bertrand competition, monopsonistic competition à la Dixit-Stiglitz 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:

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

¹⁶The direct supply function is $q_i = a + bp_i - c \sum_{j \neq i} p_j$, with $a \equiv \frac{\alpha}{\beta + \gamma(n-1)}$, $b \equiv \frac{\beta + \gamma(n-2)}{(\beta + \gamma(n-1))(\beta - \gamma)}$, $c \equiv \frac{\gamma}{(\beta + \gamma(n-1))(\beta - \gamma)}$.

¹⁷A 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.

¹⁸Appendix B.1 relaxes this assumption

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

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

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 their competitors receive represents a violation of the Stable Unit Treatment Value Assumption (SUTVA).

Treatment-Control Differences in Quantities and Prices between

This subsection derives the differences in equilibrium outcomes between treatment and control traders:

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

and

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

Δp is *increasing* in Γ : if traders are homogeneous (i.e. low Γ), control traders must charge a price closer to the treated ones to stay in the market.¹⁹ On the other hand, Δq is *decreasing* in Γ : if traders are homogeneous, the treatment traders can expand more by stealing 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 . 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, it is zero.

¹⁹If $\Gamma = 1$, 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 $\Delta p = s/2$.

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

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

The result highlights the simplification induced by the functional form we use for the supply function. The constant pass-through rate depends only on Γ and n . Intuitively, it is, decreasing in Γ and increasing in n .

3.3 Discussion

The stylized model presented in this section makes strong assumptions. We discuss some of the key ones. First, it is static. As we discuss in Section 4, 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.²⁰ Second, we rely on specific functional forms. For instance, the linear supply function that is particularly easy to work with analytically but cannot be microfounded with a discrete choice problem (see, e.g., Jaffe and Weyl, 2010 and Armstrong and Vickers, 2015). 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.²¹ 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 impact the number of players in a market, i.e., we don't allow for entry or exit.

Because of these restrictive assumptions, it is important to provide evidence in support of the choice of the model. We do this in Section 5.3. 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 B.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.

²⁰We discuss this topic further in Section 4.2.

²¹For 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.

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 in the market. 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 Empirical Results

This Section presents the empirical results of the paper. 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 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 clarifies, the subsidy affects the behavior of both treatment and control traders, as they compete in a market for the same suppliers. However, in our approach, these differences are nevertheless crucial estimation objects. By matching them to their 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 traders pay to farmers. We have data on the average cash per pound price a trader pays for a given shipment. This amount includes payments that are (potentially) made at different times (e.g., a trader advances a certain amount before harvest and pays the remaining at harvest time). We denote this variable with \tilde{p} , so to differentiate it from the *effective price*, p , we focused on in the model. Unlike the former, the latter accounts for the potential different value of payments made at different times (i.e., it weights

differently payments made at different times) and it is the object farmers ultimately care about.²²

Figure 2 displays the price results graphically. It shows weekly averages for: i) world price (right y axis);²³ 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 displays two key features. First, trader prices follow closely wholesaler prices and these move with world prices. In particular, domestic prices respond to the sharp decrease in the world price that occurred in November 2011. Second, this preliminary graphical analysis displays no obvious gap in prices that treatment and control traders pay to the farmers.

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

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

where η_z and η_t are randomization pair and week fixed effects, respectively.²⁴ We cluster standard errors at the trader level (the unit of treatment).²⁵

The term $\pi^{\tilde{p}}$ is the coefficient of interest.²⁶ Table 2, Column (1) presents a regression without week fixed effects: the coefficient is $\hat{\pi}^{\tilde{p}} = -32.5$ (se=47.16). In Column (2), we introduce week fixed effect and the coefficient becomes -6.9 (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.²⁷

One additional concern is that the treatment may induce selection in which traders make

²²We recover the treatment effect on this effective price in Section 4.2.

²³Specifically, 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.

²⁴Since pairs were matched within wholesalers, this effectively controls for the wholesaler location.

²⁵Results are similar when allowing for double clustering by trader and village (Cameron et al., 2012).

²⁶The 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.

²⁷As discussed in Section 2, six traders never visited the warehouses. Effectively, by including pair fixed effects in the regression, we drop their pairs from the analysis.

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.²⁸ 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 quite stable when including, trader controls, village controls and, finally, both set of controls. This suggests the selection concerns described above cannot drive the results.²⁹

Overall, the various specification provide persuasive 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 5.4, we use the model and our estimates to quantify the impact of the experiment on prices paid by treatment and control traders (relative to a counterfactual without the experiment).

4.1.2 Advance Payments

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

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

An observation is a farmer mentioned 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. The term π^a is the treatment effect estimator,

²⁸In 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.

²⁹In 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.

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 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 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 difference in quantities purchased between treatment and control traders. 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:

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

where an observation captures the total purchases of cocoa trader i , 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.³⁰ The results are robust when including trader controls in Column 2. Overall, this is a large impact of the treatment.³¹ Given that farmers had limited opportunities to increase production by the

³⁰We include all the eighty traders in the sample, assigning value zero to trader-week pairs with zero purchases, including for the six traders that never showed up during the experimental period). Results are similar when dropping these traders.

³¹The 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

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. In doing so, it provides estimates of the treatment effect on the *effective price*, which reflects the relative value of payments made at different times. It is important to emphasize that our framework does not model the trader choice of providing advance payments vs. raising the price. For instance, we do not have a theory of why during the experiment treatment traders are more likely than control ones to pay in advance, but not to pay a higher price at harvest time.³² As discussed in Section 3.3, 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 approach may not generate easily closed-form solutions for the differences between treatment and control and may feature multiple equilibria, thus 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. This allows us to derive the “effective price”, a measure that reflects the present value of the payment to the farmer. For this purpose, we must introduce the timing of payments in our framework. Ideally, our data would include the amounts of payments made at different times. However, we only observe whether a certain farmer receives advance payments or not. Therefore, we define the (effective) price paid by buyer i as it follows

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

where \tilde{p}_i is the total monetary amount paid by the trader, $ShareAdvances_i$ is the share of farmers to whom trader i provides advance payments. Therefore, λ is the extent to which farmers value advance payments (i.e., the rate of substitution of the indifference curve

described in Section 4.1.1.

³²We 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.

between prices and advances).³³ We estimate the value of λ using two sources of variation: cross-sectional (at baseline) and experimental.

Approach 1: Using baseline correlations

First, we look at the relation between payments (i.e. the total monetary amount paid by the trader for a given shipment) and advance payments in the cross section. 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.³⁴ Table 5, column 1 shows that moving from a village where no farmer received 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).³⁵ In terms of our model, a farmer is indifferent between a trader paying a certain price and another trader who pays 149 Leones less but provides advance payments. The result is robust to the inclusion of village level controls (column 2). These results, while based on limited baseline data, provide initial evidence consistent with the fact that farmers value advance payments and are willing to accept lower prices from traders for this service.

Approach 2: Heterogeneity in Treatment Effects

We then look at the value of advance payments using the relationship between the treatment-control differences in prices and those in advance payments. Traders respond to the treatment to by increasing the *effective* prices by a certain amount. This response can come in the form of higher payments or more frequent advance payments. The slope across the two response margins sheds light on their relative value. In other words, we estimate how much less a trader who increases his advance payments must adjust his prices.

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

³³Since 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 .

³⁴In 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. 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).

³⁵For this regression we only use data from the subset of villages in which we have pre-treatment data.

tion:

$$AdvancePayment_{f_{izv}} = \eta_z + \pi^a(\text{Treat}_i) + (\text{Treat}_i \times \mathbf{X}'_v) \boldsymbol{\pi}_v^a + \mathbf{X}'_v \boldsymbol{\beta}_v + (\text{Treat}_i \times \mathbf{X}'_i) \boldsymbol{\pi}_i^a + \mathbf{X}'_i \boldsymbol{\beta}_i + \nu_{\text{sipv}} \quad (12)$$

Where \mathbf{X}_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 compute the predicted difference in advance payment provision between treatment and controls using heterogeneity by \mathbf{X}_v and \mathbf{X}_i : $\widehat{DTC}_{iv}^a = \mathbf{X}'_v \boldsymbol{\pi}_v^a + \mathbf{X}'_i \boldsymbol{\pi}_i^a + \pi^a$. Finally, we run the following specification to test whether village-trader pairs with larger differences between treatment and control in advance payments display lower differences in prices:

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

If total payments and advance payments are substitutes (i.e., $\tilde{\lambda} > 0$), then $\pi_a^{\tilde{p}} < 0$.³⁶

Figure 4 provides initial evidence on the slope between the differences between treatment and controls along the two margins. We estimate price and advance payments differences between treatment and control 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. 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 . Any test of significance in equation (13) 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). We cluster the bootstrapping by randomization pair.³⁷

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.

³⁶We note \widehat{DTC}_{iv}^a is collinear with the vector of controls and thus cannot be included in levels in the estimating equation.

³⁷Specifically, we first estimate $\hat{\pi}_a^{\tilde{p}}$ on the full sample and generate a T-statistic, T_0 from a Wald test of the hypothesis that $\hat{\pi}_a^{\tilde{p}} = 0$. We then draw with replacement a sample of randomization pairs 1,000 times. For each draw, we predict \widehat{DTC}_{iv}^{a*} and then use it to estimate another $\hat{\pi}_a^{\tilde{p}*}$, where the star indicates the bootstrapped sample. We then generate a test statistic T^* from a Wald test of the hypothesis that $\hat{\pi}_a^{\tilde{p}*} = 0$ using the standard error from the bootstrap estimate and centering at T_0 . We hold this test statistic in memory. After 1,000 draws, a p-value is calculated from the position of (absolute value of) T_0 in distribution of test statistics T^* .

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 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 46.2 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 (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 approaches provides *causal* estimates of λ . For instance, villages with higher shares of credit provision may be systematically different and have some feature (not captured by our controls) that lower prices.³⁸ It is however reassuring that both approaches, while exploiting different variation (i.e. comparing different set of villages), lead to positive and significant value of λ .³⁹

Effective Prices

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, 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 199 when using the treatment heterogeneity (specifically, we use the specification with all the interactions, as presented in 6, col. 3). The two approaches then lead to point estimates for the treatment effect on effective price of 14.1 and 23.2, respectively. These results suggest that, in the new equilibrium, treatment traders pay an

³⁸In principle, it is also possible that some of this credit is non-interlinked. However, qualitative evidence from interaction with the traders suggest that the bulk of the credit is repaid through lower prices at harvest time.

³⁹We note that the estimates of λ derived in the second approach (214-334) are larger than the estimate from the cross sectional analysis in Table 5 (147-150). However, the baseline advance payment variable, used in the first approach, captures whether the farmer had received advances in the twelve months before the baseline, a longer time horizon than the one of the experiment (two months). Thus, the two advance payment dummies may capture different intensities of advance payments. An ideal data collection would have gathered transaction-level information on advance and harvest payments.

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, as opposed to capture 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 the trader prices respond to changes in wholesaler prices. As preliminary evidence suggesting high responsiveness, 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. Table 7 presents the results of the regression analysis. Column (1) reports a basic OLS regression of the trader price vs. the wholesaler price, with 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. These results are not directly comparable with those studies that focus on pass-through of *border* prices to final producers/consumers, such as Atkin and Donaldson (2015).⁴⁰

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=14,010). 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

⁴⁰Studies 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).

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).⁴¹

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 introduce the methodology to recover the parameters. Third, we present the estimates. Fourth, we present additional evidence in support of the framework and discuss alternative models. Fifth, we use the results to discuss the impact of the experiment (and of counterfactual experiments) on prices, quantities, and farmers margins.

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. First, we started the experiment in the middle of the harvest season. It is possible that, by that time, traders have already locked-in some farmers. Thus, the degree of competition we estimate with our intervention may have been higher if we had run the experiment ran before the harvest season. It is possible that, had the wholesalers announced the subsidy earlier, treatment traders could have accessed an even larger pool of contestable farmers and thus we may have observed a larger price response.⁴²

Second, the experiment only ran until the end of the harvest season. Traders may have behaved differently compared to a multi-season trial. For instance, treatment traders may have increased their prices less than they would have in a longer experiment. For this reason, future research should assess whether varying the duration of the experiment leads to

⁴¹Graphical 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.

⁴²As 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.

substantially different results. With this caveat, we believe that running the experiment until the end of the season is a reasonable length. The subsequent harvest season follows seven-eight months of inactivity and a lot of new trading relationships arise then.⁴³ In addition, the high pass-through rate we described in Section 4.3 suggests a high level of trader response to high-frequency price changes.

Third, another important distinction between the basic model presented in Section 3 and the experimental setting concerns the presence of non-study traders, which are about 60% of the traders operating in the market. In principle, these could be different from the study traders (control and treatment) at baseline. Importantly, the model presented in section 3 is robust to the presence of such traders. Specifically, in Appendix B, we discuss a model where only a share σ of traders is included in the study, and thus study treatments are a share $\sigma\mu$. Non-study traders have a resale price that possibly differs from that of the study traders, v . Crucially, we show that the theory treatment effects, Δp and Δq (Equations 5 and 6), and the pass-through rate ρ (Equation 7) are unchanged.⁴⁴ Therefore, the estimation approach is robust to the inclusion of non-study traders.

Finally, we discuss two additional assumptions we need to make to reconcile the theory to the available data. 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 this would bias our estimates of the advance payment treatment-control difference: traders may be less likely to extend advances to irregular supplier or, on the contrary, they may be using advances particularly to attract irregular supplier. In addition, 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). This may not be a particularly realistic assumption. Our experiment shows that advance payment provision does respond to changes in the resale price (though, 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

⁴³The fact that the experiment lasted for until the end of the season also suggests that traders had enough time to learn about their competitors subsidy, in line with the assumption of perfect information of the theory.

⁴⁴In this augmented model, Equation n is the total number of traders, i.e., study and non-study.

be a lower bound for the real value).

5.2 Methodology

Our estimation approach aims at recovering the differentiation rate, Γ , and the market size, n , which captures the number of traders effectively competing in a village. 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 a market. Most villages have a meeting point where sales occur. However, farmers may also sell 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 may only sell to traders that come at the farm gate. These arguments suggest that the definition of a market may be a complicated process. For this reason, we choose to treat the market size as an unknown, rather than assuming its value.

By matching the theoretical expressions in Equations 5 and 7 to the empirical results, we obtain a system of two equations in two unknowns, Γ , and n :

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

For a given value of the subsidy s —Leones 150,—the first equation (the experimental estimate on prices) enables us to recover the differentiation parameter Γ .⁴⁵ If we are willing to assume the value of the market size, n , the experimental results are sufficient to characterize the degree of competition, since this depends only on Γ . However, when treating market size as an unknown, the quasi-experimental estimates of the pass-through further allows us to recover n .⁴⁶

⁴⁵The estimation of the differentiation degree and of the market size is invariant to currency unit choice.

⁴⁶We 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 5.3 presents a different method that uses *percent* treatment effects in prices and quantities. The use of percent differences does not require ex ante definition of market size.

5.3 Recovering Γ and n

We solve the above system using two values of the treatment effect on effective prices, which we derived in Section 4.2, $\lambda = 149$, from the baseline cross-section, and $\lambda = 199$, from the heterogeneity in treatment effects. Solving the above system delivers the following results in terms of 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 0.18 [-.18,.59] 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 13.9 [7.6,21] with the second method.

We can always reject complete differentiation. However, due to the low precision of the estimates, we can never reject homogeneity (i.e., $\Gamma = 0$).⁴⁷ The estimated value of n is sixty to eight percent higher than 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.⁴⁸

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).

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 3.3). 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.⁴⁹ We derive theoretical expressions for the *percent* difference in prices *and quantities* between treatment and control traders and match them to their empirical counterparts (0.007 and 1.88, respectively). Using again the pass-through rate ρ as an additional moment, we recover the differentiation rate, Γ , the market size, n , and the intercept parameter, α (this latter is only identified up to a monetary unit parameter). The estimates are very similar to the previous ones. For instance,

⁴⁷When we ignore that λ is a generated regressor, we can reject $\Gamma = 0$ at $p=0.13$

⁴⁸Consistent 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.

⁴⁹Appendix C includes the details of this alternative approach.

when using our estimates of $\lambda = 199$, we obtain $\Gamma = 0.19$ and $n = 13.9$, which are very close to our previous estimates.⁵⁰ 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 set 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.

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 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). 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 to take advantage of the subsidy. First of all, it must be noted 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 could steal the suppliers of the treatment-control pair cartel). The latter form of collusion is more demanding and faces the standard enforcement problems of a cartel. We also note that collusion is *prima facie* 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. In addition, Figure 3 does not show any stark or sudden decline in quantities purchased by control traders during the intervention period (which we would expect if traders were colluding to take advantage of the subsidy). Finally, we consider standard models of monopsonistic competition (adapted from the more common monopolistic case): Dixit and Stiglitz (1977) predicts a markdown on the subsidy equal to the markdown observed in the baseline data; Ottaviano et al. (2002) would predict a difference between treatment and control traders of one-half of the subsidy value. Neither of these predictions finds support in the data.⁵¹

⁵⁰Results are similar when using $\lambda = 149$.

⁵¹Trader prices are on average 91% of the wholesaler prices. Under constant markdown case, this would

5.4 The Impact of the Experiment

Having gained confidence in the model choice, we then use the model to quantify the impact of the experiment on price, quantities, and farmer revenues (relative to a counterfactual without the experiment). We proceed in several steps.

Prices

First, through the model, we can study how the subsidy affected 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 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 45 (22) Leones in response to the experimental subsidy of 150 Leones.⁵² This confirms indeed that the difference between treatment and control, 23 Leones, is the result of a (partial) price war following the subsidy to the treatment traders.

Quantities

During the experimental period, treatment traders purchase substantially larger amounts than control ones (+190%). 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 (whose parameter depend again on Γ and n), provides a mapping from the above price impacts to the quantity impacts.

First, we find that *aggregate* quantity increases by an upper bound of 0.8%.⁵³ We can compare this result to the increase in aggregate quantity that would occur if all of the quantity treatment effects 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

imply a difference in (effective) prices of (at least) 135 Leones between treatment and control.

⁵²In setting $\mu = 0.2$, we are considering the case where non-study traders are equal to control ones. We plan to consider other cases in future versions. The results in this section are based on the value $\lambda = 214$.

⁵³The lower bound is, trivially, zero.

little way to increase their supply in response to the price changes.

Second, we can then assess the impact of the experiment on quantities purchased by treatment and control traders. Using the upper (lower) bound on the increase in aggregate quantity, we find that treatment traders increased their purchases by 112% (110%) 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% (1.8%)

Counterfactual Experiments with Different Treatment Shares

In the intervention described in this paper, only 20% of the traders received the subsidy. In this section, we compare the impact of subsidy programs targeting different treatment shares. In particular, we are interested in comparing interventions along their “return on investment”: the ratio between benefits in terms of increases in farmer revenues and their cost (the total subsidy amount). The blue curve in Figure 6 presents the results using our estimates $\Gamma = .186, n = 14$, assuming there is no impact of the experiment on aggregate quantity. We note several points. First, once shutting down the aggregate quantity impact, the return is obviously always negative: 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. This type of counterfactual exercises may be useful to inform subsidy policies.

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 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

“pass-through experiment” that targets all traders, 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 three-quarters 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. In this paper, we have developed an experimental approach where we combine 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 homogenous traders) to 1 (i.e., each trader is a local monopsonist). The empirical result suggests that provision of advance payments to the farmers is a candidate source of differentiation among traders. Second, by using quasi-experimental estimate of the pass-through rate as additional moment, we can infer the numbers of traders that compete for supply in a given village (i.e., the market size). By treating it 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: 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. 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 speaks at the difficulty of defining market size ex-ante.

Overall, consistent with the survey by Dillon and Dambro (2016), our findings suggest a competitive intermediary sector, at least for those traders who purchase from farmers. The model suggests that the experiment increased (decreased) the quantity purchased by treatment (control) traders by about 110% (27%), relative to a counterfactual without the experiment. This occurs because of market stealing rather than through an increase in aggregate quantity. 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. We hope this approach will prove useful for the many cases in which a clustered experimental design approach 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.
- 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.
- 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, 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."
- 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*.
- 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.
- 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.
- 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.
- 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.
- 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

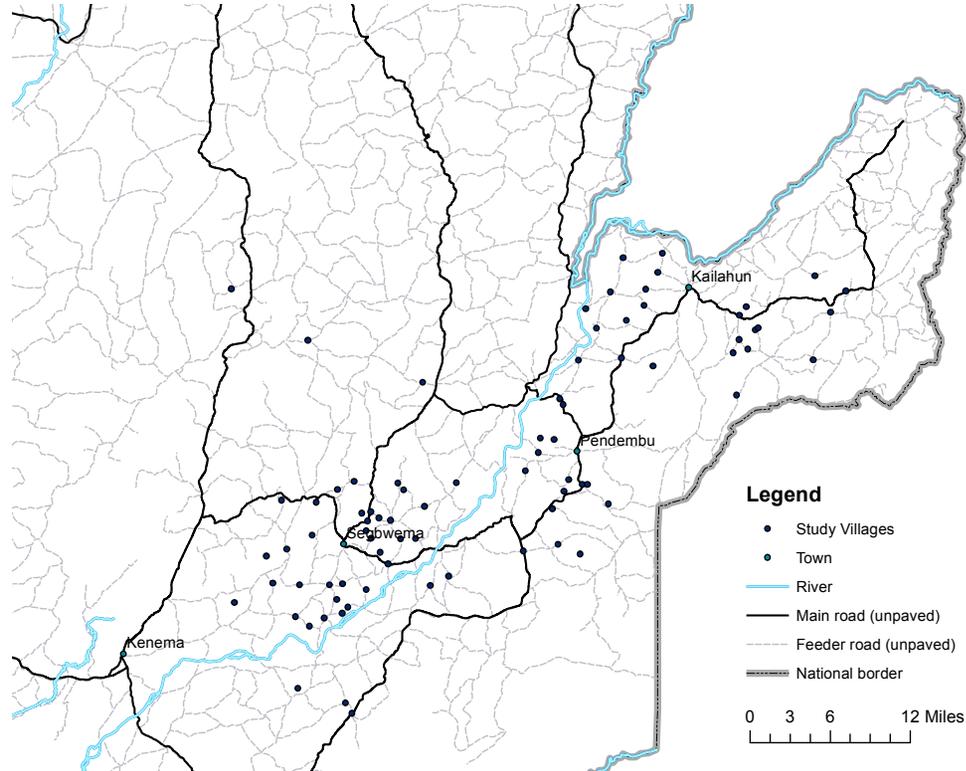
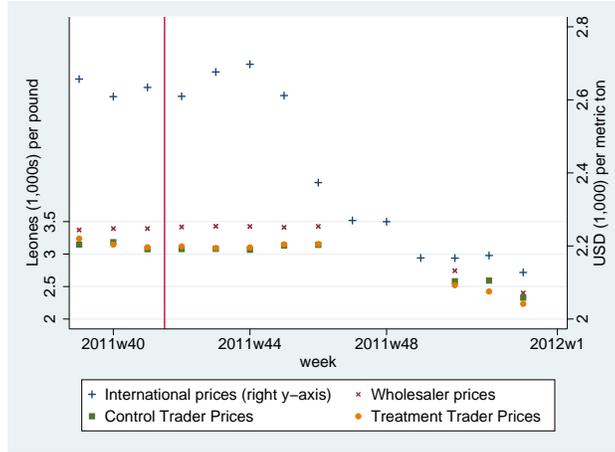
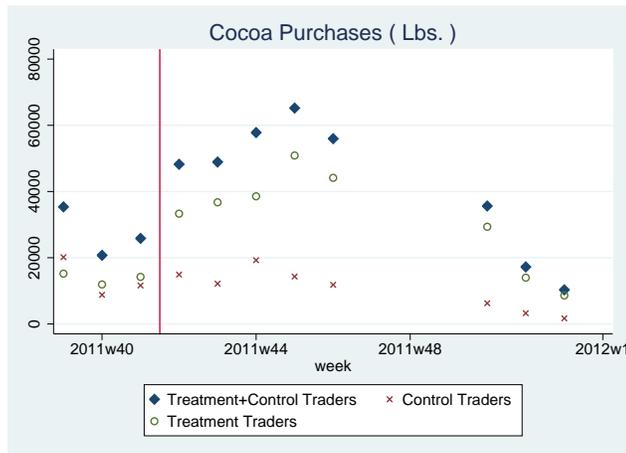


Figure 2: Cocoa Prices



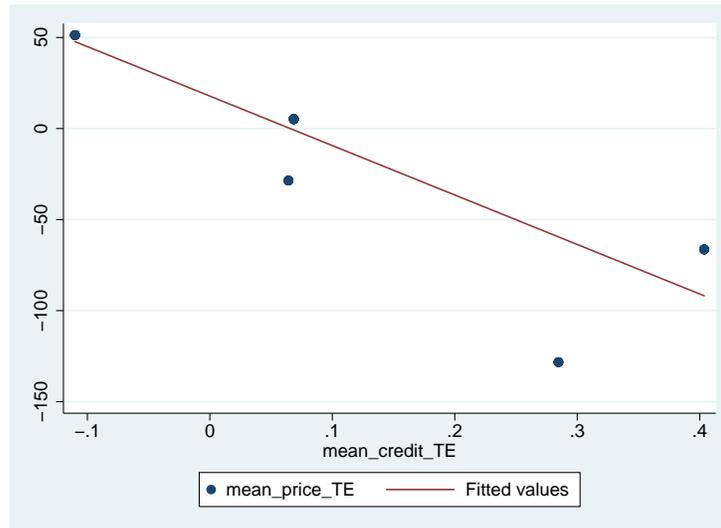
Notes: The figure presents average weekly prices for: *i*) international cocoa prices; *ii*) prices the study wholesalers pay to the traders; *iii*) prices control traders pay to farmers; *iv*) prices treatment traders pay to farmers. Wholesaler and trader prices data collection was suspended for most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

Figure 3: Purchases of Cocoa



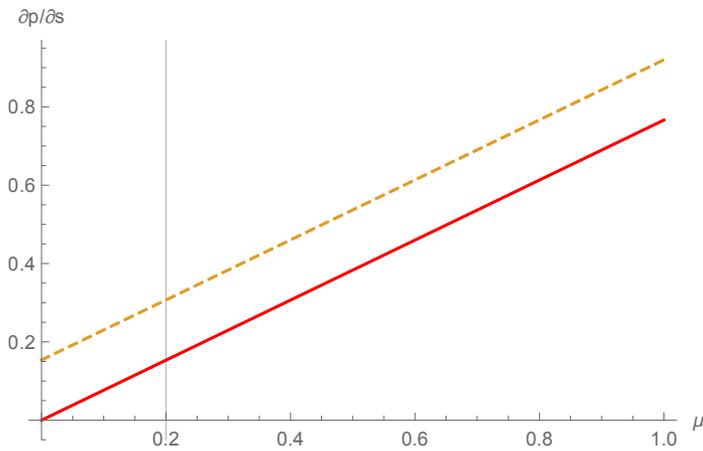
Notes: The figure shows the total amount of cocoa purchases by study (i.e., control+treatment) traders, control traders, and treatment traders. The vertical line marks the beginning of the intervention period. Data collection was suspended for most of three weeks (w47-w49).

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



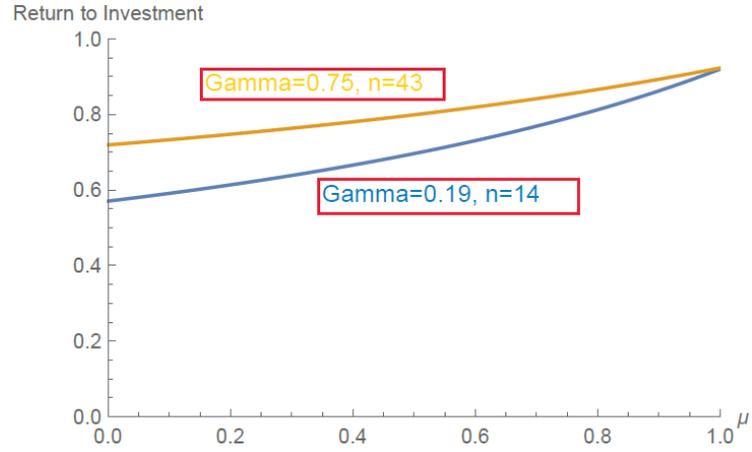
Notes: The scatter reports the the correlation across price and advance payments treatment effects, estimated separately across the five chiefdoms included in the study. The regression line has a slope of -271.

Figure 5: Simulation: Treating Different Shares of Traders



Notes: The graph shows the impact of counterfactual experiments on (“effective”) prices paid by control traders (continuous line) and treatment traders (dashed line). Specifically, it reports the increase in prices in response to a unit-subsidy 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 0.15 (0.30), respectively, relative to a scenario without the experiment. For $\mu \rightarrow 1$, the response of treatment traders tends to the pass-through rate, 0.92.

Figure 6: Counterfactual Experiments: Return on Investment



Notes: The graph reports the return to investment (y-axis) for experiments that target a share μ (x-axis) of traders with a unitary per-unit subsidy. The return to investment is defined as the ratio between the additional farmer revenues generated by the intervention and the cost of the subsidy. We conduct the simulations assuming no response in aggregate supply. Section 5.4 describes the procedure to recover these values. The blue curve describes the results using the market structure parameters estimated in the paper ($\Gamma = .186, n = 14$). The orange curve describes the results using an alternative pair ($\Gamma = .75, n = 43$) that gives the same pass-through rate, $\rho = .92$ than the pair of values estimated in the paper.

Tables

Table 1: Baseline Trader Summary Statistics

Covariate	Treatment	Control	Treatment - Control
<i>Panel A: Baseline Interview</i>			
Self-estimate bags sold in 2011	20.0	18.6	1.5 (2.23)
Age, years	38.2	36.9	1.3 (1.91)
Years selling to study wholesaler	5.7	7.3	-1.6 (0.86)*
Cement or tile floor in house $\in \{0, 1\}$	0.53	0.62	-0.1 (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 (0.39)
Number of farmers buying from	23.3	28.3	-5 (3.6)
Share of farmers given credit since March	0.72	0.62	0.04 (0.05)
<i>Panel B: Pre-treatment shipment data</i>			
Price Paid to Farmer (shipment-level)	3,136	3,137	1 (41)
Pounds sold during pre-treatment (weekly)	347	362	-5 (67)

Notes: The column “Treatment-Control” reports treatment coefficients from regressions with randomization pair fixed effects. Standard errors are clustered by trader.

Table 2: Price Responses to the Experiment

	(1)	(2)	(3)	(4)	(5)
Treatment Trader	-32.52 (47.16)	-5.47 (14.95)	-5.96 (16.94)	-12.87 (13.21)	-6.95 (15.38)
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*** (0.03)	0.14*** (0.02)	0.13*** (0.03)	0.14*** (0.02)
Control Group Mean	0.11	0.11	0.11	0.11
Trader Controls		X		X
Village Controls			X	X
Observations	1837	1837	1837	1837

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 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. 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. Share of Advance Payments: Treatment Effect Heterogeneity

	(1)	(2)	(3)
Treat* Estimated Treatment Effect on Credit	-341.79	-300.13	-198.88
<i>p-values from bootstrapped t-stats</i>	[.10]	[.15]	[.07]
Chiefdoms	X	X	X
Village Controls		X	X
Trader Controls			X
Observations	1060	1060	1060

Notes: The dependent variable is the price paid by the trader for the shipment of cocoa. Each column presents estimates of $\pi_a^{\bar{p}}$ 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*** (0.00)	0.92*** (0.00)	0.94*** (0.00)	0.93*** (0.01)	0.90*** (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.

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

B Theory Appendix

B.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.⁵⁴ 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 . By randomization, treatment is uncorrelated to firm characteristics. This highlights the key benefit of randomization even if, as we discuss extensively in the paper, one of the standard experimental assumptions, 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 homogeneity rate $\kappa\gamma$, $0 < \kappa < 1$. Therefore, the inverse demand 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

⁵⁴This is equivalent to varying producer costs in an oligopoly model.

procedure presented in the paper therefore recovers Γ and \tilde{n} .⁵⁵

B.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 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 $p_i = \alpha + \beta q_i + \gamma \sum_{j \neq i} q_j$.⁵⁶ 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).

⁵⁵The result extends to the general case of $m = 1, \dots, M$ groups of traders, with differentiation $\Gamma_M = \kappa^{m-1}\Gamma$. In this case, Δp , Δq , and ρ are as above and $\tilde{n} = \frac{n}{M} \frac{1-\kappa^M}{1-\kappa}$.

⁵⁶That 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 we use to recover the competition parameters $1 - \gamma/\beta$ and n , which we briefly discussed in Section 5.3. 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 5.2. Showing that different moments deliver similar estimates would provide support for the specific theoretical model we use.

C.1 Methodology

The main approach described in Section 5.2 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 $1 - \gamma/\beta$ 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:

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

and

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

For a given value of the subsidy s , these expressions depend on additional parameters, i.e., μ, v, α , as well as on those we aim to recover, i.e., Γ and n . We calibrate the value of μ and v . We set the former at 1/5, the share of treatment traders out of the total number of traders (study and non-study). Assigning a value to the latter requires some additional assumption. The (average) value of the wholesaler price (i.e. the price at which traders resell), is Le. 3,260. The average price at which traders purchase is Le. 2,987, 91% of the wholesaler price. However, in the model, v is the *net* resale price, net of other costs the traders may incur and that we do not observe, such as transport and storage costs. We set

$v = 3,145$, which implies a 5% markdown.⁵⁷

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 2987. The control average likelihood of advance provision is 0.11. Therefore, given $\lambda = 199$, the average control effective price is 3,009. This implies that the percent price difference between treatment and control traders during the experiment is 0.7%. The average quantity purchased by control traders is 282.5 kilograms. Thus the percent different between treatment and control traders is 188%.

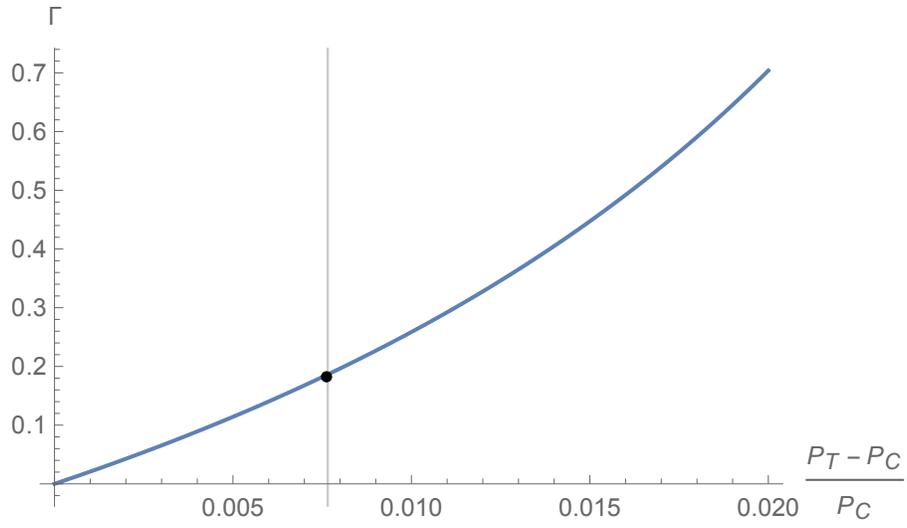
Solving the equation system with these values for $\% \Delta p$ and $\% \Delta q$, we obtain the following estimates for the three parameters of interests: $\Gamma = 0.19$, $n = 13.9$, and $\alpha = 2020$. The results for Γ and n are thus very close to the ones obtained when using the more parsimonious methodology described in the main text. We see this as evidence in support of the specific competition model chosen for the analysis.

Finally, we emphasize that the similarity of the results between the two approaches is not a mechanical result. This 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 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 effect 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 describe in Section 5.2. 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 5.3), 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).

⁵⁷Results are quite stable when using other values of v , spanning between 3,009 (the average control trader price) and 3,260 (the average wholesaler price).

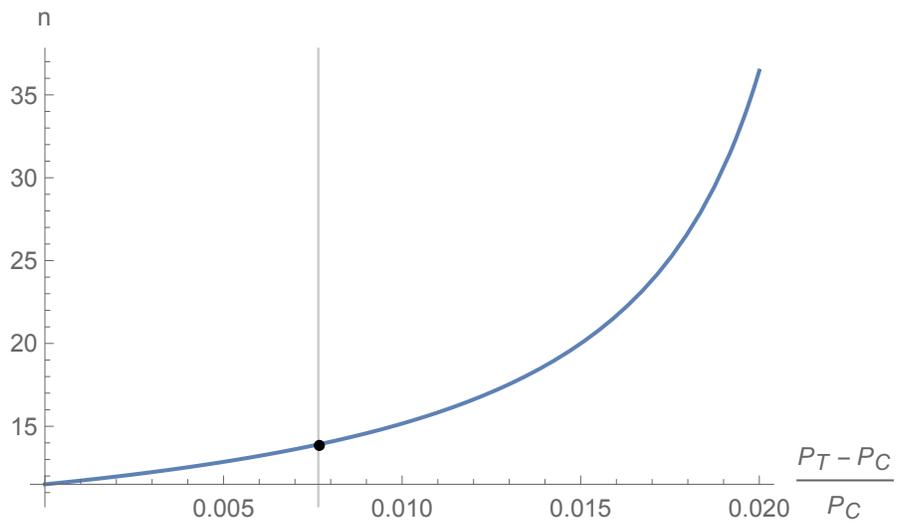
Appendix C: Figures

Figure C.1: Sensitivity of Γ to $\frac{P_T - P_C}{P_C}$



Notes:

Figure C.2: Sensitivity of n to $\frac{P_T - P_C}{P_C}$



Notes: