Online Appendix

Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya

Lauren Falcao Bergquist and Michael Dinerstein

A Appendix: Maize Value Chains and Trader and Consumer Characteristics

Figure A.1 displays the maize output market chain in western Kenya. Data for the percentage breakdown in sourcing and sale location was collected in a four-round panel survey conducted with over 300 regional traders in the area from 2013-2014 (averages displayed).

Regional traders, the subjects of this study, are responsible for large-scale aggregation, storage, and transportation. They report purchasing 50% of their maize from small and medium farmers (selling less than 5 tons), 16% from large farmers, and 33% from other traders. About half of the purchases from farmers use a local assembler or broker. Brokers are often slightly wealthier members of rural communities (and are often farmers themselves) who identify other farmers in their villages who are ready to sell. They either purchase from fellow farmers, bulk, and sell to the regional trader or, for a commission, they simply identify farmers who are willing to sell. Either way, they are small scale, often work only seasonally, and typically lack the working capital to do large-scale aggregation, long-run storage, or transport of any distance.

Traders tend to own a warehouse in a market center and either rent or own a truck which they use to purchase maize, bring it back to their warehouse for sorting, drying, and re-packaging, and then carry onward to their destination of sale. In our sample, 64% of sales take place in open-air markets in rural communities. There, 66% of traders' customers are individual households, while the rest are primarily village retailers. Traders also sell about 16% of their inventories to millers, who mill maize into flour for sale to supermarkets and other stores that serve urban consumers. They sell another 16% to other traders, who sell in other areas of Kenya or eastern Uganda. A very small portion of sales – about 2% – is sold to restaurants, schools, and other institutions. Finally, about 2% is sold to the Kenyan National Cereals and Produce Board, the former state maize marketing board that still has limited involvement in the market by purchasing, storing, and selling small reserves of maize with a goal of stabilizing prices.

Table A.1 presents summary statistics for traders in the sample. Figure A.2 displayed the average number of traders per market. The number of traders is calculated as the average number of traders present in the market during 12 weeks of the study period, as predicted by week and market fixed effects (that is, any increase in number of traders due to the entry experiment is omitted).

Table A.2 presents summary statistics for consumers served by traders in the sample. This data is drawn from a phone survey with 165 consumers randomly selected from the demand experiment sample. This survey was conducted in July and August 2016 immediately following data collection for the main experiment.

Figure A.1: **Maize value chain in study area.**

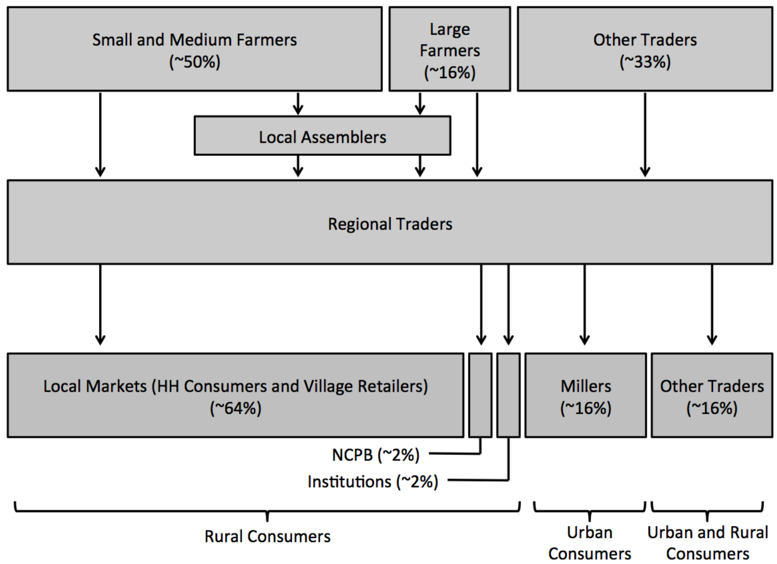

Figure A.2: **Number of traders per market**

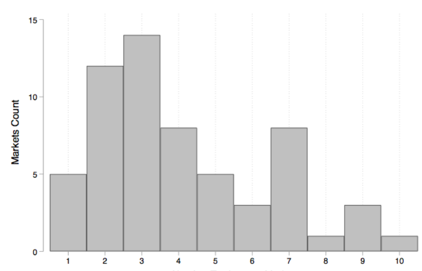

Table A.1: **Trader summary statistics.**

	Mean	Std. Dev.	Obs
Education and Business Characteristics			
Complete primary	0.78	0.42	2,728
Complete secondary	0.33	0.47	2,728
Percent corrrect Ravens	0.49	0.22	2,681
Review financial stregth monthly+	0.62	0.49	2,728
Keep written records	0.58	0.49	2,728
Any employees	0.37	0.48	2,728
Number employees	1.04	1.98	2,728
Own lorry	0.35	0.48	2,992
Market Experience			
Work in this market most weeks	0.96	0.20	2,934
New trader	0.01	0.11	2,934
Worked with all before	0.77	0.42	3,008
Know other traders well	0.68	0.47	2,549
Know other traders well or somewhat well	0.94	0.23	2,549
Single sample market trader	0.83	0.37	465
Number sample markets visited	1.29	0.77	465
Collusion Reports			
Self-report discuss price	0.38	0.49	2,549
Someone in market report discuss price	0.80	0.40	2,777
Percent traders with whom discuss price	0.77	0.28	976
Self-report agree price	0.30	0.46	2,549
Someone in market report agree price	0.72	0.45	2,777
Percent traders with whom agree price	0.77	0.28	777

Table A.2: **Consumer summary statistics.** "Number markets" is the number of markets at which the consumer typically buys maize. "Buys at least once a week" presents the percent of consumers who report buying maize at last once a week. "Search" is the percent of consumers who report approaching multiple traders before deciding from whom to buy. "Same trader" is the percent of consumers that always buy from the same trader.

	Mean	Median	SD
Number markets	1.57	1.00	1.05
Buys at least once a week	0.87		
Search	0.48		
Same trader	0.61		

A.1 External Validity

Maize is a distinctly important crop in Kenya, accounting for over a third of average gross caloric intake and about 9% of annual household expenditure Argent and Begazo, 2015. However, it is by no means unique in its market set-up, especially with regard to the physical layout of markets. The markets in which this study operates are not exclusive to maize, but rather sell a wide-variety of crops, including beans, potatoes, cabbage, tomatoes, onions, peppers, bananas, etc. For almost all crops, sellers are located immediately adjacent to each other, facilitating easy search (and potentially easy collusion). One important distinction is that while maize traders tend to exclusively sell maize, sellers of fruits and vegetables often sell several types of produce at once. Further, maize traders tend to have larger firm sizes and conduct trade across longer distances, while many produce sellers are smaller, more locally-based retail vendors.

B Appendix: Non-Nested Tests of Joint Profit Maximization and Cournot Competition

As we described in Section III, we use the profit weight model primarily to test between joint profit maximization and Cournot competition. We now turn to a non-nested test of these forms of competition, where we follow the logic of Berry and Haile (2014) and the application of Backus et al. (2019a). We refer to the results in Section VI and provide more detail here.

Our identifying logic in Section VI was that the experimental cost subsidy should be orthogonal to traders' cost type (pre-subsidy). We employ similar logic here. Under the null hypothesis that a specific model describes conduct, we can construct traders' first-order conditions. If the model is correctly specified, then the experimental cost subsidy, except for its direct effect of lowering costs by a known amount, should be orthogonal to traders' cost type and thus should be orthogonal to traders' implied marginal benefit.[1]

Specifically, we return to Equation 16 and plug in $\omega = 0$ for the null hypothesis of Cournot competition and $\omega = 1$ for the null hypothesis of joint profit maximization. For the null hypothesis of Cournot competition, we estimate:

$$P_{mw} - \Delta c_{mw} + \frac{\partial P_{mw}}{\partial q_{jmw}} q_{jmw} = \pi \Delta c_{mw} + \gamma q_{jw} + c_j + c_m + c_w + c_{jmw} \tag{B.1}$$

where we continue to use instruments because q_{jw} is endogenous but we omit the cost subsidy instruments.[2] We then test whether $\pi = 0$ where we form test statistics using our 1,000 bootstrap iterations. Our p-value on this test is 0.006 such that we reject Cournot competition.

For the null hypothesis of joint profit maximization, we estimate:

$$P_{mw} - \Delta c_{mw} + \frac{\partial P_{mw}}{\partial q_{jmw}} q_{jmw} + \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} = \pi \Delta c_{mw} + \gamma q_{jw} + c_j + c_m + c_w + c_{jmw} \tag{B.2}$$

with the same instruments. We then test whether $\pi = 0$ where we form test statistics using our 1,000 bootstrap iterations. Our p-value on this test is 0.164 such that we fail to reject joint profit maximization.

If we impose constant marginal costs and estimate via OLS, our p-values on Cournot and joint profit maximization are 0.006 and 0.180, respectively.

[1] We refer to marginal "benefit" instead of marginal revenue because under joint profit maximization the trader is not just considering his own revenues.

[2] To construct a single test, we combine the low and high cost shock treatment indicators into a single regressor: Δc_{mw}. Results are similar if we include both treatments separately.

C Appendix: Evaluating Model Assumptions

C.1 Static Model

This appendix presents the empirical basis for the decision to model a static equilibrium. Because maize is in theory a storable commodity, an alternative would be to model demand as dynamic, with prices and quantities purchased in one week affecting those bought in the next. However, empirically, consumer stockpiling is quite limited. The modal consumer purchases maize every week from her local weekly market (see Table A.2) and buys only the small amount necessary for weekly consumption (the median household consumer buys 7 kg and the median vendor buys one 90-kg bag). These weekly purchases occur against the backdrop of a 19% increase in price over the course of the lean season. If consumers were stockpiling, one would expect large purchases early in the season, when prices are low, and limited purchases later in the season, when prices are high. This is not what we observe. Related work in the region suggests that credit constraints limit households' ability to arbitrage these price fluctuations (Burke et al., 2019).

The randomized order of treatment periods allows us to go one step further and explicitly test the validity of this assumption. If inter-temporal dynamics are at play and consumers are stockpiling maize when prices drop during the pass-through experiment, one would expect a lower quantity of maize to be sold in the period following the removal of the subsidy, as consumers have stockpiled the period before. To test for this, we regress the total quantity sold in a given market-day on the previous period's treatment status (controlling for current treatment status). Column 1 of Table C.1 presents the results for the full sample. We see that having been a cost shock market in the previous 4-week block does not affect the prices, quantities sold, or number of customers in the following block. The point estimate is small in magnitude and far from statistically significant. In order to confirm that this null finding is not merely the result of low power (perhaps due to a quickly petering out stockpiling effect over the course of the 4-week block), Column 2 restricts the sample to the week immediately following the switch of treatment status, a period in which one should expect the stockpiling effect to be most concentrated. We continue to see no evidence of a stockpiling effect here (in fact, the point estimate on quantities becomes positive, though standard errors also increase substantially with this reduced sample). Given limited evidence of consumer stockpiling, we model demand as static and therefore decisions regarding prices and quantities as separable across market-days.

It is possible that the lack of effect on total quantities is the result of two competing effects canceling each other: out new customers, as they learn that the price is lower, and less demand from existing customers, as they stockpiled maize. To check for this, Columns 3 and 4 run similar specifications with the number of customers as the outcome variable. Again, we see no effects of the previous period's treatment status, suggesting this alternative explanation is not at play, and again adding confidence to the static model.

Finally, Columns 5 and 6 check for dynamic pricing, running the same specification with price as the outcome. Again we see no significant effects.

Table C.1: **Effect of Previous Treatment Status on Outcomes in Current Period.** Outcome variables as a function of previous treatment status, controlling for current treatment status. Outcome variables are log quantity sold (Columns 1 and 2), log number of transactions (Columns 3 and 4), and log price. "Cost Shock Previous" is a dummy for whether the market was in a cost shock treatment market in the previous period. Columns 1, 3, and 5 present results for the full sample. Columns 2, 4, and 6 present results for the first week of the block, when one would expect to see most concentrated dynamic effects, if existent.

	Ln Kgs	Ln Kgs	Ln Num Customers	Ln Num Customers	Ln Price	Ln Price
Cost Shock Previous	-0.0131	0.199	-0.0315	0.0539	-0.00316	-0.0108
	(0.157)	(0.213)	(0.0611)	(0.0939)	(0.00513)	(0.00657)
Mean DV	7.369	7.273	2.103	2.094	3.390	3.363
N	2191	541	2047	497	2029	495
Sample	Full Block	W1 Only	Full Block	W1 Only	Full Block	W1 Only
Market FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes

It is also worth noting that storage is also quite limited among traders. Burke et al. (2019) find that "in a panel survey of local traders, we record data on the timing of their marketing activities and storage behavior, but find little evidence of long-run storage." In that data, collected with traders from the same region, only 31% of traders report doing any storage. Those that do on average store about 30% of the bags they buy and only for extremely short periods of time (on average, 4 days, among those that store). Only 1.2% of traders store maize for more than a week. For the rest of the traders, our general supply model accommodates such arbitrage as the cost function covers all sales in the week.

C.2 Product Differentiation

Staple food commodities are often pointed to as the textbook example of a homogenous goods. However, we take seriously the concern that this assumption could be wrong and that there could be quality differences across sellers, which would result in product differentiation. We therefore collect detailed quality estimates. Note that the use of grain standards in Kenya is restricted to the most formal settings of large millers and the National Cereals and Produce Board. Regional traders typically do not know the official grade of their maize, and consumers do not use grades to describe or evaluate quality. Instead, traders and consumers assess quality of maize based on several readily observable characteristics: coloration, grain size, grain intactness, presence of foreign matter, and presence of weevil infestations. Therefore, we measure quality according to the these standards, which are those relevant to the market actors in question. Enumerators were trained to grade quality on a scale from 1 (lowest quality) to 4 (highest quality) according to the following rubric, which was developed with the guidance of several traders in the pilot: 4=Excellent [no pest, no foreign matter, no broken grain, no discoloration, sizable grain]; 3=Good [barely infested, <5% foreign matter (e.g., maize cobs, dust, sand etc.), <5% broken grain, <5% discolored]; 2=Fair [infested, 5%-25% foreign matter, 5%-25% broken grain, 5%-25% discolored]; 1=Poor [infested, >25% foreign matter, >25% broken grain, >25% discolored].[3]

There is no variation in quality offered by a single trader to his customers in the same market-day. In fact, it is common for traders to mix bags they have purchased of different quality prior to arrival at the market with the explicit goal of offering a uniform quality level.[4] We therefore collect only one measurement of quality for each trader in each market-day. Across traders in the same market day we observe little variation in quality, as measured on a scale of 1-4 (97% of all maize receiving a rating of 2 or 3). Moreover, as shown in Column 1 of Table C.2, prices are not statistically different across the (limited) variation seen in quality.

The other salient dimension on which products might be differentiated is the availability of credit (while not strictly a dimension of the physical product, the ability to buy on credit is dimension of the transaction). However, credit does not appear to be a salient factor in these primarily "cash-and-carry" spot markets; over 95% of transactions are conducted in cash. That said, it may be that the *availability* of credit matters to a minority of customers; when asked how customers decide on which trader from whom to buy, 34% cite the availability of credit when needed, so it does appear that a slightly larger percent of customers value the possibility of obtaining a line of credit in periods when they are in need. Moreover, while we do see small price differences for purchases on credit, this relationship disappears when controlling for other features of the transaction.[5]

Locational differences, combined with search costs, could also be a basis for product differentiation. However, within a given market, search costs for consumers are negligible in this setting, as traders sell in trucks parked immediately next to each other or in stores located immediately adjacent to each other

Reflective of this limited variation in product characteristics (e.g. quality, credit, etc.), we see little variation in prices. The coefficient of variation in prices offered by the same trader, same day is 3.1%, while the coefficient of variation in the average price of traders in the same market, same day is 5.1%.

Therefore, the weight of evidence appears to suggest that maize sold in these markets is a relatively homogenous good.

Finally, it is worth noting that even if differentiation were driving the low pass-through we estimate, much of the paper's framework and conclusions would still be relevant. We would still interpret the low pass-through, given the same demand estimates, as evidence that the

[3]No formal tools were used to measure precise percentages; rather, enumerators were trained to take a handful of maize in their palm and count the kernels that matched each description. While this involves some imprecision, it is nearly identical to the process by which consumers judge quality – that is, by feel, sight, etc. – and therefore captures well the information available to consumers, which is the pertinent metric. Enumerator training on grading included practice evaluating the quality level of real samples of maize.

[4]Incentives to maintain a uniform average quality could be driven by consumer preferences or by a desire to not deviate from the average quality offered by other traders.

[5]Unexpectedly, the relationship between credit and price seen in Column 2 is negative, but this may be driven by omitted variables such as transaction size and consumer identity. After controlling for these factors in Column 3, there is no significant difference in price charged for credit transactions (and the coefficient is now sensibly positive, albeit very small in magnitude).

trader is capturing most of the change in surplus from the cost decrease.[6] Our demand model includes customer fixed effects, which could represent customer heterogeneous preferences or some additional utility the customer derives from buying from that trader, especially because the demand experiment occurs after the customer has chosen a trader. The main adjustment would be altering the supply side to determine how heterogeneous consumers sort to different traders. Thus, while the source of market power would be different, it would still likely lead to traders capturing most surplus gains from cost decreases.

Table C.2: **Product Differentiation.** Data drawn from trader price surveys, broken out by transaction (there are almost 40,000 transactions observed in the full dataset). Market-day fixed effects are employed to compare difference in transaction characteristics only within the same market-day. Quality is ranked on a scale from 1(=lowest quality) to 4(=highest quality). Credit is a dummy for whether the transaction was conducted on credit. Other controls refer to the size of the transaction and the identity of the customer (household vs. village retailer). All standard errors are clustered at the trader x date level.

	(1) Ln Price	(2) Ln Price	(3) Ln Price
Quality (1-4, 4=best)	0.000450		0.00156
	(0.00212)		(0.00180)
Credit		-0.0177	-0.000767
		(0.00273)	(0.00276)
Mean Dep Var	3.366	3.366	3.366
N	39598	39667	39598
Market-day FE	Yes	Yes	Yes
Other Controls	No	No	Yes

C.3 Price Discrimination

Empirically, we see little variation in the price that a given trader offers his customers through the day; the intra-cluster correlation of these prices is 0.9. While there is no official posted price to ensure that prices are equivalent across customers, negotiations between traders and customers occur in public (often in front of the trader's truck or store, where other customers are typically lined up to purchase). This likely limits traders' ability to engage in dramatic price discrimination. However, traders may be able to engage in some small and imperfect price discrimination using tools such as bulk quantity discounts, as documented in recent work by Attanasio and Pastorino (2015).[7] To explore whether there is evidence of such nonlinear pricing schemes in our setting, we utilize transaction-level data (totaling 39,667 transactions) and explore the covariance of price and quantity of maize sold by the same

[6]Note that any costs the trader incurred to supply a differentiated product are already incorporated via heterogeneous marginal cost or are fixed costs that do not change in response to the cost decrease.

[7]Attanasio and Pastorino (2015) find that sellers of food staples in Mexico are able to exert market power to discriminate across customers with different levels of willingness (and ability) to pay. Sellers in their setting offer nonlinear pricing schemes using bulk discounts.

trader to his customers in a given market-day. Figure C.1 presents this relationship, plotting a kernel-weighted local polynomial regression of log price on log quantity, both demeaned by trader x market-day fixed effects. While the relationship is relatively flat in the middle of the distribution, we see that customers at the lower end of the quantity distribution are paying more per kg, while those at the higher end are paying less per kg. The 95% confidence interval area, delineated in grey, suggests that these bulk discounts are particularly prominent at very large quantities. The effect sizes are relatively small, with the bulk of overall variation of price lying within a band of about +/-1%; however, they do suggest that traders possess some limited ability to use nonlinear pricing to price discriminate. Note that any ability to price discriminate is prima facie evidence of market power.

Figure C.1: **Quantity discounts**. Within trader x market-day residuals of transaction-level log price/kg and quantity/kg. N=39,667. Grey area represents the 95% confidence interval.

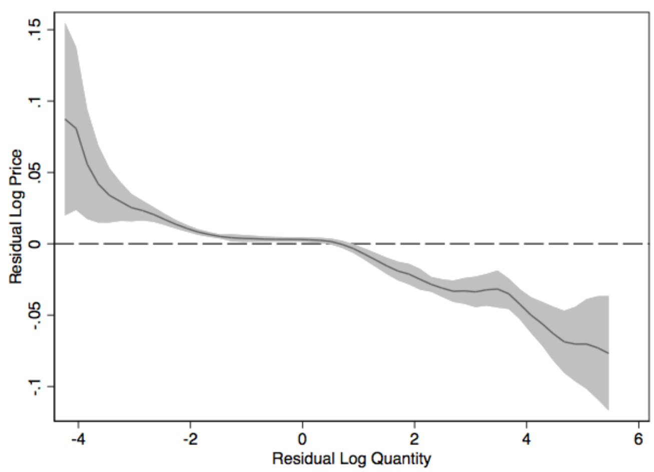

D Appendix: Constant Marginal Costs

A key assumption of the model underpinning the simple model is that of constant marginal costs. In this appendix, we present direct empirical evidence, beyond the more general model, supporting this assumption.

This evidence suggests that the assumption of constant marginal cost is in fact a fairly good fit for the empirical setting. Agricultural intermediation is an industry for which the majority of variable costs – the purchase price of the inventory, the cost of casual laborers' time for loading and off-loading, etc. – appear to be fairly constant with respect to quantities. While there may be a discontinuous increase in marginal cost when capacity constraints are hit (for example, if a trader sells more than the capacity of his truck and would need to bring a second truck to sell an additional bag), empirically this constraint is rarely binding, as only 7% of traders in the sample sell out of the full amount of maize they have brought to the market that day. Consistent with this, a detailed investigation of trader expenses across three countries finds that traders appear to face fairly constant costs across these settings (Fafchamps et al., 2005).

This is concordant with the estimates from our general model. Our estimate of γ, the marginal cost slope, is 0.0006 Ksh/kg, and 0 is well within a fairly tight 95% confidence interval of $(-0.0006, 0.0016)$. This point estimate is small, implying that a 1 standard deviation increase in weekly (in-sample) quantity sold (2300 kgs) corresponds to a cost increase of just 1.73 Ksh/kg. This is small compared to the heterogeneity in trader-market-week marginal cost intercepts – marginal cost for the first kg – where we estimate a standard deviation of 10.84 Ksh/kg. Thus, given this auxiliary cost data plus the structural model estimates, we consider constant marginal costs as a reasonable approximation of the empirical setting in which this experiment takes place.

E Appendix: Sample Selection and Experimental Schedule

The sample of markets in this study is drawn from six counties in Western Kenya. These counties encompass most of the (Kenyan) area within a 50km radius from the town of Bungoma, Kenya, the site of the research hub for this study. A listing exercise was conducted with the Director of Trade in each county to get a comprehensive list of all markets in the county. We excluded markets that were reported to not have any maize traders typically present. These represent some of the smallest rural markets, which have only maize retailers, who in turn purchase their maize from traders in larger markets. Major urban markets in the town centers were also excluded since the primary focus of this study is on the rural markets frequented by rural consumers.[8]

The exercise yielded 154 potential markets for inclusion. From this sample, 60 markets were selected in the following stratified manner: 40 markets were selected from within a radius of 50 km of Bungoma town and 20 markets were selected from outside this radius.[9] We administered a pre-experiment survey to this group of 60 selected markets in which we verified information provided by the Director of Trade and recorded the number of traders typically in the market.[10] In a large number of these markets, it was found that the information provided by the Director of Trade was inaccurate.[11] Markets that were deemed ineligible upon visit were then replaced with market from their same stratum.[12] Newly selected markets were then visited in an identical verification exercise. This process was continued until 60 markets had been selected for inclusion in the sample.

Figure E.1 presents the experimental schedule. The 60 markets in our sample are randomly assigned one of six possible schedules, in order to yield randomized ordering of treatment statuses. There are therefore 10 markets in each schedule. This allows the inclusion of market and week fixed effects in every analysis. There is therefore a total of 720 market days in our sample, clustered into 180 market x four-week block cluster (standard errors in all specifications are clustered at this market x four-week block level). The demand experiment is run in a quarter of the markets during each week break in between each treatment status. Each market therefore receives the demand experiment once.

[8]These markets represented only 2% of the total markets listed.

[9]The 40 markets within 50km of Bungoma were selected randomly. This randomization was stratified to include 25 markets from which we had valuable historical data from pilot work, while the remaining 15 markets were new to the sample. The 20 markets located more than 50km from Bungoma were selected according to a non-random algorithm in order to minimize confounding effects due to spillovers and get a larger geographic distribution of markets. For each market, the distance to the nearest market in the pool (the 40 selected markets within 50km of Bungoma as well as any remaining markets in this outer circle pool) was calculated and then the market with the shortest distance was dropped.

[10]Each trader present in the market during this verification exercise was asked "How many maize traders are typically present in this market on an average market day from March to July?" Answers were averaged across all traders to yield a single measure of the number of traders typically present in the market.

[11]The most common issue being that the market was so small as to not have any traders.

[12]That is, markets from the first stratum forming the area within 50 km of Bungoma were replaced with another randomly selected market from this stratum. Markets from the outer stratum of 20 markets were replaced with the next further market, according to the algorithm determining selection in this stratum.

Figure E.1: **Experimental schedule.**

<table>
<tr><th></th><th>Schedule 1</th><th>Schedule 2</th><th>Schedule 3</th><th>Schedule 4</th><th>Schedule 5</th><th>Schedule 6</th></tr>
<tr><td>Week 1</td><td colspan="6">Demand Experiment in 1/4 of markets</td></tr>
<tr><td>Week 2</td><td rowspan="4">Pass Through</td><td rowspan="4">Control</td><td rowspan="4">Entry</td><td rowspan="4">Pass Through</td><td rowspan="4">Control</td><td rowspan="4">Entry</td></tr>
<tr><td>Week 3</td></tr>
<tr><td>Week 4</td></tr>
<tr><td>Week 5</td></tr>
<tr><td>Week 6</td><td colspan="6">Demand Experiment in 1/4 of markets</td></tr>
<tr><td>Week 7</td><td rowspan="4">Entry</td><td rowspan="4">Pass Through</td><td rowspan="4">Control</td><td rowspan="4">Control</td><td rowspan="4">Entry</td><td rowspan="4">Pass Through</td></tr>
<tr><td>Week 8</td></tr>
<tr><td>Week 9</td></tr>
<tr><td>Week 10</td></tr>
<tr><td>Week 11</td><td colspan="6">Demand Experiment in 1/4 of markets</td></tr>
<tr><td>Week 12</td><td rowspan="4">Control</td><td rowspan="4">Entry</td><td rowspan="4">Pass Through</td><td rowspan="4">Entry</td><td rowspan="4">Pass Through</td><td rowspan="4">Control</td></tr>
<tr><td>Week 13</td></tr>
<tr><td>Week 14</td></tr>
<tr><td>Week 15</td></tr>
<tr><td>Week 16</td><td colspan="6">Demand Experiment in 1/4 of markets</td></tr>
</table>

F Appendix: Estimation Details

In this appendix we provide estimation details for the empirical models in the main text. We start by specifying all model equations before providing estimation details for each component. Our general model is:

$$q_{imt}(P_{imt}) = \begin{cases} \left(\frac{a_i - P_{imt}}{b_i}\right)^{\frac{1}{\delta}} \eta_{imt} & if \;\; P_{imt} \leq a_i \\ 0 & if \;\; P_{imt} > a_i \end{cases} \tag{F.1}$$

$$a_i \sim N(\mu_a, \sigma_a^2) \tag{F.2}$$

$$Q_{mw}(P_{mw}) = \sum_{i \in \mathcal{I}_{mw}} q_{imw}(P_{mw}) \tag{F.3}$$

(F.4)

$$P_{mw} - \Delta c_{mw} + \frac{\partial P_{mw}}{\partial q_{jmw}} q_{jmw} = -\omega(Entry, Contacts) \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} + \gamma q_{jw} + c_j + c_m + c_w + c_{jmw}$$

$$\omega(Entry, Contacts) = \begin{cases} \omega_n & if \;\; \text{No Entry} \\ \omega_e^{with} & if \;\; \text{Entry by Trader with Contacts} \\ \omega_e^{without} & if \;\; \text{Entry by Trader without Contacts} \end{cases} \tag{F.5}$$

$$Entry_{jmw} = \begin{cases} 0 & if \;\; \pi^V_{jmw}(MC^0_{jmw}, \omega) < FC_{jmw} - EntrySubsidy_{jmw} \\ 1 & if \;\; \pi^V_{jmw}(MC^0_{jmw}, \omega) \geq FC_{jmw} - EntrySubsidy_{jmw} \end{cases} \tag{F.6}$$

$$\begin{pmatrix} MC^0_{jmw} \\ FC_{jmw} \end{pmatrix} \sim logN\left(\begin{pmatrix} \mu_{MC} \\ \mu_{FC} \end{pmatrix}, \begin{pmatrix} \sigma^2_{MC} & \rho_{MCFC}\sigma_{MC}\sigma_{FC} \\ \rho_{MCFC}\sigma_{MC}\sigma_{FC} & \sigma^2_{FC} \end{pmatrix} \right). \tag{F.7}$$

Equations F.1 and F.3 are household and market demand, respectively. Equation F.4 describes the supply side and nests Cournot and collusion. Equation F.6 determines trader entry and Equation F.5 lets the form of competition change with entry (ω_e indicates entry, ω_n indicates no entry). Equations F.2 and F.7 impose distributional assumptions on some of the unobserved heterogeneity. This gives us the following parameters:

Parameter	Description
δ	Curvature of individual demand
μ_a	Mean of demand intercept
σ_a	Standard deviation of demand intercept
b_i	Heterogeneity in price coefficient
γ	Marginal cost slope
c_j	Trader-specific marginal cost
c_m	Market-specific marginal cost
c_w	Week-specific marginal cost
ω_n	Profit weight if no entry (baseline equilibrium)
ω_e^{with}	Profit weight if entry by connected entrant
$\omega_e^{without}$	Profit weight if entry by unconnected entrant
μ_{MC}	Mean marginal cost intercept
μ_{FC}	Mean fixed cost
σ_{MC}	Standard deviation of marginal cost intercepts
σ_{FC}	Standard deviation of fixed costs
ρ_{MCFC}	Correlation between marginal cost intercept and fixed cost

Because we estimate model separately, we list each set of moments below depending on the equation.

F.1 Estimating Demand

We take the log of Equation F.1 and then take first differences within consumer (where the two time periods are before the offered subsidy and after). This gives us the following equation to estimate:

$$\text{(F.8)} \quad log(q_{im1}) - log(q_{im0}) = \frac{1}{\delta}\left(log(a_i - P_{im1}) - log(a_i - P_{im0})\right) + \left(log(\eta_{im1}) - log(\eta_{im0})\right)$$

where $d_i \equiv P_{im1} - P_{im0}$ is the subsidy (or discount) amount. d_i takes on 10 values, but because consumers randomized to the zero subsidy treatment were not offered a chance to change transacted quantity, we drop this group such that the remaining sample has 9 different subsidy values.

Our first set of moments is $E(\mathbb{1}\{d_i = d\}\left(log(\eta_{im1}) - log(\eta_{im0})\right))$ for the 9 different values of d. Our second set of moments draws from the control and cost shock markets. Let t_{mw} be the transaction rate in market m in week w and q_{mw} be the mean quantity (kgs) per transaction.[13] Let t^c, t^{low}, and t^{high} be the sample mean transaction rates for control, low cost shock, and high cost shock market-weeks, respectively, with analogous notation for mean

[13] We construct t_{mw} from the number of observed transactions, dividing by the maximum number of transactions observed in that same market in a single week during the 12-week experimental period. Our results are robust to increasing this denominator by at least a factor of 2.

quantity per transaction. Further, let P^c, P^{low}, and P^{high} be the sample mean prices in these three treatment arms.[14] Then we construct the following six moments:

- $E(\mathbb{1}\{a_i > P^c\} - t^c) = 0$
- $E(\mathbb{1}\{a_i > P^{low}\} - t^{low}) = 0$
- $E(\mathbb{1}\{a_i > P^{high}\} - t^{high}) = 0$
- $E(\mathbb{1}\{a_i > P^c\}\left(\frac{a_i - P^c}{b_i}\right)^{\frac{1}{\delta}}\eta_{imt} - q^c) = 0$
- $E(\mathbb{1}\{a_i > P^{low}\}\left(\frac{a_i - P^{low}}{b_i}\right)^{\frac{1}{\delta}}\eta_{imt} - q^{low}) = 0$
- $E(\mathbb{1}\{a_i > P^{high}\}\left(\frac{a_i - P^{high}}{b_i}\right)^{\frac{1}{\delta}}\eta_{imt} - q^{high}) = 0$

The sample values of these moments are:

Table F.1: **Quantity Moments from Cost Experiment.** Transaction rate is the number of transactions divided by the maximum number of transactions observed in any week for a given market. Kgs/Transaction is the average kgs per transaction in any week for a given market.

	(1) Transaction Rate	(2) Kgs/Transaction
Low Cost Reduction Treatment	0.267	-32.63
	(0.0258)	(11.26)
High Cost Reduction Treatment	0.284	-14.68
	(0.0547)	(18.38)
Constant	0.241	72.25
	(0.0576)	(39.17)
Mean Dep Var	0.589	98.27

We use these 15 moments in estimating δ, μ_a, and σ_a via method of simulated moments where we simulate from the normal distribution of a_i. We use a three-step (iterated) procedure with an estimated optimal weighting matrix and use an analytical gradient to speed up computation. We do not estimate b_i directly, as it drops out of our first difference specification. But the distribution of b_i, or more specifically, $b_i^{-\frac{1}{\delta}}\eta_{imt}$, is necessary for calculating the model predicted moments and for subsequent analysis. Given estimates of δ, μ_a, and σ_a, we estimate the distribution of $b_i^{-\frac{1}{\delta}}\eta_{imt}$ with the following procedure:

1. Draw a transaction (quantity-price pair) from the set of transactions in control, low cost shock, and high cost shock market-weeks.
2. Draw a_i from $N(\hat{\mu_a}, \hat{\sigma_a^2})$, which we have assumed is independent across consumers.

[14] If we use the full vector of market-week prices in sample we get similar results.

3. Compute $b_i^{-\frac{1}{\delta}}\eta_{imt}$ to rationalize the chosen quantity-price pair.

4. Repeat until have sampled all transactions from the data.

To construct 95% confidence intervals, we run 1,000 bootstrap iterations where we draw two iterations in each sample. First, we resample (with replacement) a set of market-blocks from the control, low cost shock, and high cost shock markets, to recompute the last 6 sample moments. Then we resample (with replacement) the set of consumers from the demand experiment.

With estimates of individual demand, we can estimate market demand. The key object to estimate is $\mathcal{I}_{mw}$, the number of consumers per market-week.[15] We use our demand estimates to simulate consumers and predict their demand, given the observed market-week price. We draw consumers until in aggregate their predicted demand matches Q_{mw}, the observed quantity transacted in the market-week.

We then estimate $\frac{\partial Q_{mw}}{\partial P_{mw}}$ for each market-week. Given the functional form for demand,

$$\frac{\partial Q_{mw}}{\partial P_{mw}} = \sum_{i\in\mathcal{I}_{mw}} \frac{-1}{\delta(a_i - P_{mw})}\left(\frac{a_i - P^{mw}}{b_i}\right)^{\frac{1}{\delta}}\eta_{imt} = \sum_{i\in\mathcal{I}_{mw}} \frac{-1}{\delta(a_i - P_{mw})}q_{imw} \tag{F.9}$$

where in the last step we plug in transacted quantities from the data.

The functional form also lets us calculate consumer surplus:

$$CS_{mw} = \sum_{i\in\mathcal{I}_{mw}} \frac{\delta}{1+\delta}(a_i - P_{mw})q_{imw} \tag{F.10}$$

F.2 Estimating Supply, without Entry

On the supply side, we estimate:

$$P_{mw} - \Delta c_{mw} + \frac{\partial \hat{P}_{mw}}{\partial q_{jmw}}q_{jmw} = -\omega_n \frac{\partial \hat{P}_{mw}}{\partial q_{jmw}}\sum_{k\neq j} q_{kmw} + \gamma q_{jw} + c_j + c_m + c_w + c_{jmw} \tag{F.11}$$

with two-stage least squares. The only differences between F.4 and F.11 is here we plug in estimated inverse demand derivatives and given there is no entry in our main supply model, we specify $\omega(Entry, Contacts) = \omega_n$.

In addition to the sets of trader, market, and week fixed effects, we have 8 excluded instruments:

- an indicator for whether the market is in low cost shock treatment (1 moment)

[15] One option would be to use the maximum number of transactions observed in that same market in a single week during the 12-week experimental period, as we did above in constructing our moments. But because the number of consumers per market-week is often not too big, our estimate may be far from the observed data if, say, consumers in a market-week draw particularly high values of a_i. To generate more precise estimates, we therefore incorporate observed transacted quantities in each market-week.

- an indicator for whether the market is in high cost shock treatment (1 moment)
- the fraction of a trader's markets in each treatment group (low cost shock, high cost shock, entry) (3 moments)
- indicators for whether the trader has a low, medium, or high subsidy to enter *a different* market (3 moments)

Our moment conditions are $E(IV * c_{jmw} = 0|j, m, w)$ for each of the 8 instruments. The first two instruments are orthogonal to cost type, even unconditionally, based on the experimental randomization. The last 6 require conditioning on the trader.

We estimate standard errors and confidence intervals with 1,000 bootstrap iterations where we resample (with replacement) at the market-block level, which was the level of experimental randomization. The sampling is the same as the demand estimates, and we bootstrap demand and supply jointly.

After testing whether $\omega_n = 0$ or 1, and finding that $\omega_n = 1$, we impose $\omega_n = 1$ and re-estimate Equation F.11. We use these re-estimated cost parameters when calculating markups and profits and in all subsequent analysis.

F.3 Estimating Supply, with Entry

We then estimate whether entry changes how traders compete. We start by estimating via two-stage least squares a pooled entry effect: ω_e:

(F.12)

$$P_{mw} - \Delta c_{mw} + \frac{\partial P_{mw}}{\partial q_{jmw}} q_{jmw} + (1 - Entry_{mw}) \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} = -\omega_e Entry_{mw} \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} + \gamma q_{jw} + c_j + c_m + c_w + c_{jmw}$$

where we imposed $\omega = 1$ if no entry occurs. We add an extra instrument: whether the market-week is in the entry experiment. We construct our regressors using all of the traders, but we only include non-entrants as observations in estimation.[16]

We then examine heterogeneity in entry effects based on entrants with and without connections to traders in the market. Using two-stage least squares, we estimate:

(F.13)

$$P_{mw} - \Delta c_{mw} + \frac{\partial P_{mw}}{\partial q_{jmw}} q_{jmw} + (1 - Entry_{mw}) \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} = -\omega_e^{with} EntryWith_{mw} \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} - \omega_e^{without} EntryWithout_{mw} \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} + \gamma q_{jw} + c_j + c_m + c_w + c_{jmw}$$

[16] We re-estimate the cost parameters – for use only in testing ω_e – rather than using the estimates from above, as the sample of traders is somewhat different now that entry treatment markets are in the sample.

where "With" indicates entry by a trader with connections and "Without" indicates entry by a trader without connections. We add one more instrument: whether the potential entrant who received the high entry subsidy has connections in the market.

For both models, we estimate standard errors and confidence intervals with 1,000 bootstrap iterations.

F.4 Estimating Entry

Let $MC^0_{jmw} = c_j + c_m + c_w + c_{jmw}$ be trader j's marginal cost intercept in market m in week w. We estimate Equation F.6 using method of simulated moments where we draw marginal and fixed costs according to Equation F.7 (evaluated at candidate parameter values), calculate variable profits if the potential entrant were to enter, and then use the model to determine whether the potential entrant would actually enter.

Estimating variable profits involves finding a new market equilibrium in quantity choices. For candidate quantity choices, we estimate the market price using estimated inverse demand, and we solve for an equilibrium where all traders choosing positive quantities have their first-order conditions hold. We estimated incumbents' costs above, which allows us to search for a new equilibrium.

Let $Takeup^{Low}$, $Takeup^{Med}$, and $Takeup^{High}$ be the sample entry take-up rates for the potential entrants receiving low, medium, and high subsidies, respectively. Let $EntryMC^{Low}$, $EntryMC^{Med}$, and $EntryMC^{High}$ be the estimated marginal cost intercepts time entry for the potential entrants receiving low, medium, and high subsidies, respectively.[17] Let $PredTakeup^{Low}$, $PredTakeup^{Med}$, $PredTakeup^{High}$, $PredEntryMC^{Low}$, $PredEntryMC^{Med}$, and $PredEntryMC^{High}$ be the model predictions for the same objects.

We specify 6 moments:

- $E(PredTakeup^{Low} - Takeup^{Low}) = 0$
- $E(PredTakeup^{Med} - Takeup^{Med}) = 0$
- $E(PredTakeup^{High} - Takeup^{High}) = 0$
- $E(PredEntryMC^{Low} - EntryMC^{Low}) = 0$
- $E(PredEntryMC^{Med} - EntryMC^{Med}) = 0$
- $E(PredEntryMC^{High} - EntryMC^{High}) = 0$

We estimate via method of simulated moments, where we estimate in two steps to use an estimated optimal weighting matrix. To ease the computational burden of needing to solve for a new equilibrium for every market-week for every set of candidate parameters, we use importance sampling for simulating the marginal cost distribution.[18] We use $\mu = 3$ and

[17] As we estimate marginal cost intercepts above, we can only estimate them for actual entrants. Thus, by interacting with entry, we only need estimates for entrants' marginal costs.

[18] Fixed cost draws do not alter the quantity-setting equilibrium, conditional on entry, so we simulate in two steps. First we simulate marginal costs using importance sampling. Then we simulate from the fixed cost distribution, conditional on each marginal cost draw.

$\sigma = 1$ as parameters for the importance sampling distribution. We simulate 25 marginal cost draws per potential entrant-market-week. For certain starting values we run into estimates of a degenerate distribution, a common problem in importance sampling that introduces considerable simulation error (Ackerberg, 2009). We impose a lower bound on the parameters of 0.25 and avoid degeneracy.[19]

We estimate standard errors with 1,000 bootstrap iterations.

F.5 Counterfactuals

In counterfactuals, we alter the form of competition, and recompute equilibria in quantity choices in each market-week. For candidate quantity choices, we estimate the market price using estimated inverse demand, and we solve for an equilibrium where all traders choosing positive quantities have their first-order conditions hold. In the counterfactual with exit, if a trader is making negative total profits, we remove him from the market and compute a new equilibrium. We iterate until no remaining trader makes negative total profits. If multiple traders make negative total profits in a given equilibrium, we remove the trader making the largest losses and recompute the equilibrium.

[19] Results are similar for other lower bounds.

G Appendix: Heterogeneity in Entry Effects

This appendix explores pre-specified dimensions of heterogeneity in the effect of entry on competition. Specifically, we pre-specified three dimension of heterogeneity: whether the entrant has contacts in the market, whether the entrant is large (above median profits), and whether the entrant's ethnicity matches that of the majority of traders in the market.

We already explored how the form of competition varies with whether the entrant has contacts in the market (Table 6). Here we report the results from the other two pre-specified sources of heterogeneity: entrant size and entrant ethnicity. We see that our estimates to too imprecise to make strong conclusions on heterogeneity based on either dimension.

Table G.1: **Effect of Entry on Competition.** The tables shows separate profit weights depending on (1) whether the entrant has above or below median profits and (2) whether the entrant's ethnicity corresponds to the majority ethnicity among traders in the market. The second and third columns show the bounds of the 95% confidence interval, calculated with 1,000 bootstrap iterations.

Group	**Parameter Estimate**	**95% CI LB**	**95% CI UB**
Heterogeneous by profits			
ω^{above}_e Entrants with above median profits	0.52	-0.28	1.38
ω^{below}_e Entrants with below median profits	0.98	0.20	1.68
Heterogeneous by ethnicity			
ω^{maj}_e Entrants in ethnic majority	0.58	-0.09	1.23
ω^{min}_e Entrants in ethnic minority	1.29	-3.81	12.66

H Appendix: Quantity Effects in Cost Shock Experiment

We have explored several candidate explanations of the large demand response to the cost shock experiment. First, we consider the question of data quality. One concern might be that if busy markets led enumerators to potentially miss some transactions, and subsidized traders had an incentive to have their transactions recorded, this could generate a large extensive margin elasticity. We believe our data is robust to this concern. Enumerators were stationed at the trader's location of sale, visually monitoring each transaction. While we cannot rule out that enumerators may have missed some transactions during busy market days, we intentionally allocated more enumerators to markets that were particularly busy, with the explicit aim of minimizing this.

A second possibility is storage, which might mean that *market* residual demand is substantially more elastic than *consumption* demand. However, as addressed in Appendix C.1, there is limited evidence of meaningful storage in our setting.

A third possibility is that consumers are substituting across markets. While our survey indicates that most consumers are "captured" by the local market, it is possible that some consumers substitute from nearby markets outside our sample (control markets in our study were intentionally spread out sufficiently to avoid spillover concerns) and that this contributes to the increase in quantities observed in treatment markets. We find somewhat greater evidence in support of this explanation. We take the census of all markets in the six counties in which our study was run (this is the sample frame from which we initially randomly sampled markets for inclusion in our experiment). We identify, for each market in our sample, the number of "potential substitute" markets by counting the number of other markets occurring on the same day of the week within a variety of radii surrounding the market. We then test whether the increase in quantities observed in the experiment is concentrated in markets with a greater number of neighboring markets, as one would expect if substitution were at play. We also look specifically at the effect of having large neighbor markets, as these markets may contain a larger number of consumers who may switch into the study market and prompt a large quantity response.

Table H.1: **Quantity Effects by Number of Neighbors.** "Kgs" is the total kg sold in any week for a given market. "Cost Change" ("CC") is the change in cost (Ksh/kg) from treatment: -2.22 in the low cost shock treatment and -4.44 in the high cost treatment. "Num Neigh Xkm" is the number of markets within X km that have the same market-day and "Num Large Neigh Xkm" is the number of such markets that are large according to the market census.

	(1) Kgs	(2) Kgs	(3) Kgs	(4) Kgs	(5) Kgs	(6) Kgs	(7) Kgs
Cost Change	-846.8	-615.2	-681.8	-670.4	-676.3	-665.0	-663.2
	(160.5)	(187.8)	(127.1)	(139.6)	(120.2)	(131.0)	(119.2)
CC x Num Neigh 10km		-349.6					
		(301.0)					
CC x Num Large Neigh 10km			-1143.6				
			(453.6)				
CC x Num Neigh 5km				-953.6			
				(552.6)			
CC x Num Large Neigh 5km					-2527.0		
					(557.5)		
CC x Num Neigh 3km						-1687.8	
						(865.3)	
CC x Num Large Neigh 3km							-3850.6
							(132.4)
Mean Dep Var	4627.9	4627.9	4627.9	4627.9	4627.9	4627.9	4627.9
N	474	474	474	474	474	474	474
Market FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Results indicate that the quantity response seems to be related to the number of potential substitute neighboring markets, and that this is driven mostly by having *large* neighbors. We also see that the coefficients increase in magnitude as we focus on geographically closer neighbors. We see very similar results when we look at the number of transactions:

Table H.2: **Transactions Effects by Market's Neighbors.** "Trans" is the total number of transactions in any week for a given market. "Cost Change" ("CC") is the change in cost (Ksh/kg) from treatment: -2.22 in the low cost shock treatment and -4.44 in the high cost treatment. "Num Neigh Xkm" is the number of markets within X km that have the same market-day and "Num Large Neigh Xkm" is the number of such markets that are large according to the market census.

	(1) Trans	(2) Trans	(3) Trans	(4) Trans	(5) Trans	(6) Trans	(7) Trans
Cost Change	-10.80	-11.50	-9.907	-10.16	-9.953	-10.32	-10.13
	(1.838)	(2.571)	(1.949)	(2.132)	(1.877)	(1.998)	(1.885)
CC x Num Neigh 10km		1.051					
		(2.162)					
CC x Num Large Neigh 10km			-6.194				
			(3.256)				
CC x Num Neigh 5km				-3.473			
				(4.168)			
CC x Num Large Neigh 5km					-12.57		
					(3.994)		
CC x Num Neigh 3km						-4.465	
						(4.806)	
CC x Num Large Neigh 3km							-14.16
							(6.485)
Mean Dep Var	58.73	58.73	58.73	58.73	58.73	58.73	58.73
N	474	474	474	474	474	474	474
Market FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

In terms of magnitude, how much of the total quantity response is explained by this effect? We do a simple back-of-the-envelope calculation based on these regression results. Taking the 5-km radius results (a reasonable travel distance in this setting), we see in Table H.1 Col. 5 that markets with no neighbors see a 676 kg increase in quantity transacted per one unit change in the cost reduction, and this effect increases by 2527 kg for each same-day large neighbor. The mean market has 0.068 same-day large neighbors within 5 km. Thus, we can calculate the total treatment effect as $-676 - 2527 * 0.068 = -848$ and the same-day large neighbor component as $-2527 * 0.068 = -172$, or 20% of the total treatment effect.[20]. Of course, this specification by no means captures all forms of substitution, so we view these back-of-the-envelope calculations as merely suggestive that substitution from large, same-day neighboring markets explains a relevant portion of the demand elasticity, but not necessarily

[20] If we focus on the all neighboring markets specification in Col. 4, we see again that the effects are primarily driven by the large markets. The mean market has 0.169 same-day neighbors (of any size) within 5 km. If we multiply the regression coefficient in Col. 4 by the mean number of neighbors within 5 km, we get -161; if we multiply the regression coefficient in Col. 5 by the mean number of large neighbors within 5 km, we get -172. Thus, the effect of seems driven entirely by large neighbors.

all of it.

To the extent that this substitution explains at least some of the large demand response, our inference on the form of conduct appears robust. The market residual demand elasticity is what we need to identify the degree of competition in the market. We are estimating traders' optimal response to cost shocks in their market, and this is dictated by the market residual demand curve they face. As long as the estimated demand curve reflects the change in a trader's quantity transacted from lowering price, regardless of where these consumers are coming from, we can model traders' incentives and thus use their choices to infer the degree of competition.

However, cross-market substitution could potentially create spillovers that contaminate the control group and bias our estimates of pass-through. We think this is unlikely for several reasons. First, our study sample includes only a random 60 markets from among the full census of 225 markets in the study counties, and therefore most of the markets from which additional customers are drawn are markets that are not in our study, rather than control markets. Moreover, study markets were intentionally chosen to be spread out from each other to mitigate this issue. Indeed, just 22% of neighboring markets (within a 5km radius) are in the study sample. This is particularly the case for large markets, from which the above evidence suggests most of these substituting customers are drawn. Of the 38 large markets in our study area, only two are within a 5km radius of any other same-day study market. Second, we find that while the quantity effects depend in part on neighboring same-day markets, there is minimal evidence that pass-through rates do (see Table H.3). [21]

[21] We can also directly control for any potential spillover effects in a Miguel-Kremer (2004) style specification, and we continue to find pass-through rates of around 22%.

Table H.3: **Price Effects by Market's Neighbors.** "Cost Change" ("CC") is the change in cost (Ksh/kg) from treatment: -2.22 in the low cost shock treatment and -4.44 in the high cost treatment. "Num Neigh Xkm" is the number of markets within X km that have the same market-day and "Num Large Neigh Xkm" is the number of such markets that are large according to the market census.

	(1) Price	(2) Price	(3) Price	(4) Price	(5) Price	(6) Price	(7) Price
Cost Change	0.224 (0.0434)	0.225 (0.0499)	0.210 (0.0468)	0.221 (0.0511)	0.220 (0.0451)	0.224 (0.0481)	0.229 (0.0459)
CC x Num Neigh 10km		-0.00167 (0.0362)					
CC x Num Large Neigh 10km			0.0910 (0.0627)				
CC x Num Neigh 5km				0.0128 (0.0633)			
CC x Num Large Neigh 5km					0.0508 (0.0931)		
CC x Num Neigh 3km						-0.00157 (0.0687)	
CC x Num Large Neigh 3km							-0.0954 (0.0526)
Mean Dep Var	28.92	28.92	28.92	28.92	28.92	28.92	28.92
N	1860	1860	1860	1860	1860	1860	1860
Market FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Third, as a robustness check, we verify that these large markets with same-day neighbors, which might be affected by neighboring markets' treatment status, do not affect our conclusions about the form of competition by dropping them from our analysis. In the simple model, the cost experiment only affects the inference about competition through the pass-through rate. In Table H.4 we see that the estimated pass-through rate is nearly identical when we use the full sample of markets, drop all large markets, or drop the two large markets that are within the 5km radius of any other study market with a common market-day ("donors"). We conduct a similar exercise for the general model. Recall that using the full sample, we have a point estimate on ω, the profit weight, of 1.07. If we drop the two large markets that are within the 5km radius of any other same-day study markets, we estimate a profit weight of 1.01. Our conclusions are therefore robust to restricting our sample to either markets that are unlikely to draw consumers away from other markets (Table H.3)) or to markets that are unlikely to lose consumers to other markets.

Table H.4: **Price Effects by Market's Same-Day Neighbors and Market Type.** "Cost Change" ("CC") is the change in cost (Ksh/kg) from treatment: -2.22 in the low cost shock treatment and -4.44 in the high cost treatment. The "No Large" sample excludes markets that are large according to the market census. "No Donors" excludes the two large markets that have a same market-day neighbor within 5km.

	(1) Price	(2) Price	(3) Price
Cost Change	0.224	0.231	0.232
	(0.0434)	(0.0517)	(0.0465)
Mean Dep Var	28.92	29.10	29.00
N	1860	1255	1743
Market FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Sample	All	No Large	No Donors

I Appendix: Entry Effects in Cost Shock Experiment

In using the cost shock experiment to infer how traders compete, a maintained assumption is that the set of traders does not endogenously change in response to the cost shock. However, it is possible that the shock introduces entry of new traders to take advantage of the subsidy. We investigate this possibility by estimating the effect of the cost shock on the number of entrants to the market:

Table I.1: **Entry Effects.** "Cost Reduction Market" is a dummy for treatment status in the cost-reduction experiment. "Number Entrants" is the number of traders present in the market on that day who had never worked in that market before.

	(1) Number Entrants
Cost Reduction Market	0.0338 (0.0175)
Mean Dep Var	0.0484
N	1860
Market FE	Yes
Week FE	Yes

We are able to classify entrants, even in the cost shock experiment, because the first time we encountered a trader during our study period, our survey asked the him about his past experience working in that market. We find a very small – albeit marginally significant – effect on the number of entrants. Note that – to the extent that we do see entry in the cost-shock experiment – this will mean we are less likely to infer collusion. To see this, imagine two scenarios. First, suppose that entry occurs but that conduct is unchanged. Having more traders in the market would weakly lower prices. We would thus expect the experimental pass-through estimate to be larger in magnitude than the estimate we would get were there no entry. This would bias us away from collusion. Second, suppose that entry occurs and that it changes conduct. This would likely be a change toward more competition, which would be a second reason for prices to fall. Again, the experiment pass-through estimate would be larger in magnitude than the estimate we would get were there no entry. As a final robustness check, we re-estimate our general model, but drop the traders that entered. This re-defines the moment condition to be that the cost shock is orthogonal to incumbents' cost residuals, and avoids any potential endogeneity of particularly low- or high-cost traders endogenously entering and causing our original moment condition to be violated. When we estimate this model, we get an estimate of ω of 1.09 (vs. 1.07 in the baseline specification). Thus, we do not believe that this small entry effect changes our conclusions.

J Appendix: Additional Details on the Experimental Design

This appendix provides additional detail on the experimental design of the three experiments.

J.1 Experiment 1: Trader Cost Shock Experiment

When introducing the subsidy, enumerators first asked the trader to describe some of the major costs that he faced in his business (traders in control market days were also asked these questions, to avoid confounding treatment with any priming effects). The subsidy was then framed as a reduction of these costs. At no point were traders told that the purpose of the subsidy was to see how much would be passed on to the prices they set for customers; rather they were told the research was interested generally in how "reductions in cost affect your business." To reduce the chance that traders viewed the subsidy as a gift, we explicitly stated the following in the script that described the subsidy: "We would wish to offer [X] Ksh for every bag you will be able to sell in this market today. This amount of [X] Ksh will offset the costs you incurred to get these bags to the market today. Remember, [X] Ksh is directed towards cutting costs and is NOT a personal gift or a promotion."

To ensure trader comprehension, enumerators then guided traders through a check for understanding. This included asking the trader if he understood the rules (99.8% of traders reported they did). To confirm this understanding, the enumerator then asked the trader to explain back to enumerators, in his own words, the meaning of the cost-reduction subsidy and the rules by which it operated. Using this method, enumerators confirmed that 99.6% of traders understood the rules (the most crucial of which were the size of the subsidy and the fact that it was tied to the number of bags sold). A similar procedure was used to ensure that traders understood the duration of the treatment, with traders describing back to enumerators in their own words how long the subsidy would run. 96.8% of traders were reported to have understood the duration of the intervention "well," 3.0% "somewhat well," and only 0.2% "not well." Finally, and most importantly, payments were set-up in a two-step procedure each day, with the explicit goal of building trader trust. Traders received the first payment early in the day (following the first hour of sales) and the second payment at the end of the day, following the completion of the day's sales. Payments were sent to traders' phone via M-Pesa (mobile money), so they received them in real time. This was designed (1) to build trader trust that they would, in fact, receive payment and (2) to ensure that traders experienced how the amount of the subsidy was calculated as a function of the number of bags sold, which is the key feature that traders must understand to align their incentives with the theory tested in the experiment. These features ensured that traders understood and trusted the structure of payments and their implied incentives for pricing.

We did not inform consumers about the subsidy. Rather, we left it to traders to determine how much information to share with consumers, as they would naturally.

We took several steps to prevent fraudulent sales. Most importantly, enumerators were stationed with each trader, monitoring each transaction, and were therefore able to observe cash and maize exchange hands. A random subset of customers was selected for additional monitoring, including questions on customer identity, the purpose of the purchase, etc. Managers stationed in the market were required to be present for this interview and approve any transactions above a certain size as subsidy-eligible.

The subsidy ran for four weeks. Exploratory interviews conducted with traders prior to the implementation of the experiment suggested that this matched well the duration at which traders face naturally occurring cost shocks, such that both traders and consumers would find this duration of shock to be commonplace and would respond naturally. Traders were specifically asked whether consumers would react badly if they lowered their prices during the subsidy and raised them in the future, following the removal of the subsidy. Traders stated this was not a major concern, as the duration of the shock matched other cost shocks during the season, such that would not find any resulting price shock to be unusual.

J.2 Experiment 2: Demand Experiment

In the demand experiment, customers were first allowed to approach traders and negotiate a price and quantity in a natural way before being approached by an enumerator to invite them to the demand experiment. If the customer consented, a random discount amount was drawn (using a randomization feature within SurveyCTO) and the customer was told that the price he had previously received from the trader would be reduced by that amount. The customer was then invited to select a new quantity he would like to purchase in light of this new price. Consumers were permitted to return home to collect any additional cash required to make their desired purchase, if needed. The price discount was given to the customer in the form of a mobile money or a cash transfer, and the customer paid the trader the originally negotiated price.

Table J.1 presents balance by baseline price and quantity demanded, by subsidy level. The p-value from an overall F-test is presented at the bottom of the table and demonstrates that we cannot reject that all coefficients are zero.

Table J.1: **Demand Balance.** Pre-subsidy pice is presented in Column 1, while pre-subsidy quantity is presented in Column 2. The p-value from an overall F-test is presented at the bottom.

	Price	Quantity
Subsidy Level 2	0.23	-27.58
	(0.53)	(17.36)
Subsidy Level 3	-1.16	0.57
	(0.54)	(17.68)
Subsidy Level 4	-0.32	-25.61
	(0.55)	(17.89)
Subsidy Level 5	-0.45	-19.92
	(0.55)	(17.89)
Subsidy Level 6	-0.30	-13.23
	(0.53)	(17.39)
Subsidy Level 7	-0.63	-3.30
	(0.53)	(17.28)
Subsidy Level 8	-0.90	-21.31
	(0.55)	(17.93)
Subsidy Level 9	-0.70	-0.00
	(0.52)	(17.08)
Subsidy Level 10	-0.57	-10.65
	(0.55)	(17.93)
Mean Dep Var	32.06	78.04
N	1361	1361
F-Test	.31	.62

Several checks were put into place to prevent consumers from making multiple visits until they received a larger subsidy. First, enumerators were stationed in the market for the full day and were trained to identify such returning customers. Second, prior to revealing the discount amount, enumerators recorded the name and phone number of the consumer, which could be used to check for previous subsidy assignment. Consumers would therefore have to give false names and phone numbers if they were to revisit with the goal of receiving a larger subsidy. Because the subsidy was delivered via mobile money to the phone number listed, this discouraged reporting of false phone numbers.

J.3 Experiment 3: Entry Experiment

In the entry experiment, traders who had never before worked in the treated market were offered subsidies to enter and attempt to sell there. The pool of traders eligible to receive the entry offers was drawn from the sample of traders interviewed in pilot work (traders from markets in the same region in Kenya) and the universe of all traders found during the market census activity. Small traders who did not own or regularly rent trucks were then excluded from the pool as pilot work showed that these traders categorically did not take up the offer. A phone survey was conducted with the remaining 187 traders to determine markets in which they had ever worked. For each of the 60 sample markets, we then identified the set of eligible traders who (1) had never before worked in that market and (2) did not work in other study markets that occur on the same day of the week in order to avoid inducing exit in our sample. The median market had 37 eligible traders, the minimum had 28, and the maximum had 56. From each of these sets, we then randomly selected the three traders who would receive the entry offers.

Because we did not want to overwhelm a single trader with too many offers, we only offered each trader one offer per 4-week block. Because this has cascading effects for the set of eligible traders for each market, we randomize the order in which markets were assigned traders from the remaining pool. In the first block, a few traders asked to be removed from the study (due to lack of interest in the subsidy and therefore unwillingness to answer surveys). When these traders were scheduled to receive an offer in a subsequent block, they were then replaced and the offer was given to a new, unassigned trader from the same pool.