# Pass-through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya

Lauren Falcao Bergquist*

September 8, 2017

Find the most updated version **here**

**Abstract**

African agricultural markets are characterized by low revenues for farmers and high food prices for consumers. Many have worried that this wedge is partially driven by imperfect competition among intermediaries. This paper provides experimental evidence from Kenya on intermediary market structure. Experimentally elicited parameters governing cost pass-through and demand curvature are used to calibrate a structural model of market competition. Estimates reveal a high degree of intermediary market power, with large implied losses to consumer welfare and market efficiency. Exogenously induced firm entry has negligible effects on prices and competitiveness parameters, implying that marginal entry does not meaningfully enhance competition.

**JEL Classifications:** D22, D43, F12, L13, L81, O13, Q13

*Becker Friedman Institute, University of Chicago, 1126 E. 59th Street, SHFE 214, Chicago, IL 60637. Email: lfbergquist@uchicago.edu. For their helpful comments, I thank Edward Miguel, Marshall Burke, Benjamin Faber, Willa Friedman, Benjamin Handel, Kaushik Krishnan, Nicolas Li, Jeremy Magruder, Jamie McCasland, Craig McIntosh, Odyssia Ng, Katalin Springel, Michael Walker and seminar audiences at Boston University, Columbia, Cornell (Dyson), International Food Policy Research Institute, Microsoft Research, NYU, Pennsylvania State University, SITE (Stanford), UC Berkeley, UC Santa Cruz, University of Chicago, University of Michigan, University of Tennessee, World Bank Development Research Group, Yale, and WGAPE (Stanford). I also thank Meshack Odhiambo Okello, Ben Wekesa, Deborrah Muthoki Wambua, and Innovations for Poverty Action for excellent research assistance in the field. This research was supported by funding from the NSF Graduate Research Fellowship and the research initiative "Private Enterprise Development in Low-Income Countries" (MRG004_3854), a program funded jointly by the Centre for Economic Policy Research (CEPR) and the Department for International Development (DFID). The views expressed are not necessarily those of CEPR or DFID. The design for this experiment was registered with the AEA RCT Registry (RCT ID: **AEARCTR-0000682**). All errors are my own.

The 1980s and 1990s saw a wave of liberalization sweep across African agricultural markets as part of broad structural adjustment plans. Inherent in the promise of these reforms was the presumption that a competitive private sector would emerge to take advantage of newly created arbitrage opportunities, with agricultural traders efficiently moving crops from surplus to deficit regions, and from harvest to lean seasons. However, recent empirical estimates suggest that agricultural markets remain poorly integrated, with prices varying widely across regions and seasons (Moser, Barret and Minten, 2009; Burke, Bergquist and Miguel, 2016). High transaction costs contribute to this limited market integration. Transport costs in Africa are the highest in the world (Teravaninthorn and Raballand, 2009); also prevalent are harder-to-measure costs associated with search (Aker, 2010), contractual risk (Startz, 2017), and price uncertainty (Dillon and Dambro, 2016).

However, much less is known about the degree of competition among intermediary traders in agricultural markets in developing countries. Whether traders are exerting market power matters for policymaking: if intermediaries are operating in a competitive environment in which price gaps are purely due to high transactions costs, then policies that reduce these transaction costs – road improvements, preferential terms for business expansion loans, and trade intelligence systems for broadcasting prices to traders, for example – would yield savings that traders will pass on to farmers (in the form of higher prices) and consumers (in the form of lower prices). On the other hand, if traders are exercising market power, gains from policies that reduce traders' operating costs may not be fully passed on to farmers and consumers; instead, the bulk of these benefits may be captured by intermediaries. To meaningfully improve farmer and consumer welfare in this environment, policies may need to explicitly target enhanced competition among intermediaries.

In this paper, I present some of the first experimental evidence on the market structure in which African agricultural traders operate. To this end, I implement three randomized control trials that are tightly linked to a structural model of market competition. In particular, I use new empirical evidence on the extent of pass-through, the shape of demand, and the effects of entry on market prices in order to contribute to our understanding of the welfare implications of imperfect competition in this setting.

In the first experiment, I exogenously reduce traders' marginal costs by offering to all traders in a market a substantial, month-long subsidy per kg sold. I then observe how much of this reduction in costs is passed through to the price offered to consumers. I find that traders pass through only 22% of this reduction in costs to customers, substantially less than the 100% pass-through predicted in a simple perfectly competitive model.

Nonetheless, the pass-through rate is insufficient to characterize imperfect competition as the curvature of demand could produce lower pass-through rates, holding behavior of intermediaries constant. For example, the observed rate of pass-through could be consistent with a Cournot competitive market structure with highly concave demand or with a perfectly collusive market structure with moderately concave demand. In order to distinguish between the roles played by intermediary conduct and consumer demand curvature, which is necessary to quantify the severity of the deviation from perfect competition, I run a second experiment to estimate the curvature of demand. In this experiment, I offer consumers random reductions in price spanning a range of counterfactual pass-through rates and measure the resulting quantities purchased. I use these results to structurally estimate a highly flexible parametric demand function.

To quantify the competitiveness of agricultural intermediaries, I use these experimental estimates of pass-through and demand curvature to calibrate a structural model motivated by the framework proposed in Weyl and Fabinger (2013) and Atkin and Donaldson (2015). Results indicate that the degree of competition is low. In fact, the estimated parameter governing competitiveness is statistically indistinguishable from that representing a perfectly collusive model in which traders form agreements (perhaps tacitly) about prices and act as a single profit-maximizing monopolist in the market. I can rule out more familiar forms of competition, such as Cournot competition and perfect competition, with 90% confidence.

Using these estimates for welfare analysis, I find that imperfect competition in these agricultural markets reduces total variable surplus by 14.6%. Of the remaining surplus, intermediaries capture 79% percent while consumers enjoy a mere 21%. Counterfactual simulations suggest large increases to consumer welfare from greater competition. These gains are driven in large part by a transfer of surplus from intermediaries to consumers, though they are augmented by a reduction in deadweight loss.

My third experiment tests whether policies that incentivize market entry can decrease market power and promote competition. The literature has struggled to empirically identify the impact of entry, which is an endogenous response to market conditions. I generate exogenous entry by incentivizing traders to enter randomly selected markets for the first time. The experiment results in an additional 0.6 traders per market-day on average, a 13% increase over the mean market size (and 20% over the median).

I use the model to solve for the predicted price changes resulting from entry under various counterfactuals for entrant behavior. Given estimated demand parameters, counterfactual simulations predict that the entry generated by the experiment should decrease prices by

8% if the entrant competes, 4% if conduct among traders remains unchanged, and 0% if the entrant colludes. This compares to the precisely estimated observed drop of 0.5% in the experiment, which is strong suggestive evidence of collusion. Structural estimates which jointly estimate demand and the change in the competitiveness parameter also find parameter estimates consistent with entrants colluding with incumbents. These results suggest that collusive agreements among intermediaries are flexible and can readily accommodate new entrants. Results from this paper therefore cast doubt on the power of entry by a small number of new traders to dramatically improve market competition in this setting.

This paper is one of the first to experimentally test the competitiveness of rural agricultural markets directly. Previous attempts to measure competition have mainly relied on observational methods. Observational studies have typically found high rates of pass-through across major markets (Rashid and Minot, 2010), though these high transmission rates may not extend beyond major urban markets (Moser, Barret and Minten, 2009). Moreover, interpretation of this observational evidence is confounded by common shocks such as shared harvest times and reverse flows across seasons. One exception to this primarily observational literature is a recent paper by Casaburi and Reed (2016), which studies the effect of an experimental subsidy per unit purchased offered to cocoa traders in Sierra Leone. They find small pass-through in terms of price, but larger pass-through in credit, suggesting the importance of interlinked relationships in their context (a feature not relevant in the Kenyan maize markets I study, in which over 95% of transactions are conducted in cash).[1]

Another set of papers attempts to directly measure traders' profits in order to draw inference about the size of rents and degree of competition. These have generally found that average trader profits are high, though subject to large variability, leaving a question mark on whether these large returns represent rents or risk premia (Dillon and Dambro, 2016). Moreover, these direct measures are subject to severe measurement error in the face of difficult-to-quantify search, own labor, and fixed costs (Fafchamps, Gabre-Madhin and Minten, 2005), as well as the sensitivity of directly asking about profits in an environment in which traders are often labeled as exploitative.[2]

[1] Because their subsidy is offered only to a subset of traders in the market, Casaburi and Reed (2016) must ultimately rely on observational estimates of pass-through to measure the degree of competition, as their experimental estimates appear to be affected by within-market spillovers. Further, in the absence of evidence on the shape of farmer supply, they are forced to make strong linearity assumptions. Because the curvature of the market facing traders (farmer supply in their case, consumer demand in mine) is crucial to interpreting the pass-through rate, I experimentally estimate this curvature.

[2] A lack of record keeping exacerbates these challenges of direct measurement. Only 58% of traders in this sample keep any written records; and, among this group, most records are fairly rudimentary.

Finally, a set of papers has applied experimental methods to the somewhat related question of the impact of offering price information to farmers on their ability to extract better prices from traders. While most studies find null results (Fafchamps and Minten, 2012; Mitra et al., 2015),[3] it is unclear if this suggests traders are already offering competitive prices given their transport costs or whether farmers are simply unable to utilize this information to improve their bargaining position. There is therefore a paucity of causally identified evidence on trader competitiveness (Dillon and Dambro, 2016) despite a growing interest in the role these intermediaries play in determining the allocation of gains from trade (Antras and Costinot, 2011; Bardhan, Mookherjee and Tsumagari, 2013).[4]

Theoretically, this paper is closely related to the framework developed in Atkin and Donaldson (2015). They use the pass-through rate of cost shocks to non-agricultural goods in Nigeria and Ethiopia to adjust for variable mark-ups in trade cost estimates. In this paper, I experimentally estimate pass-through in order to apply this method to an agricultural setting in which ubiquitous domestic production and consumption make it difficult to cleanly trace price shocks from distinct points of origin.[5] Further, I extend this work by identifying how key model parameters governing competition respond to entry. More generally, I add to the recent literature using experimentally estimated parameters to understand market structure in developing countries (Keniston, 2011).

This paper proceeds as follows: Section 1 describes maize markets in Kenya. Section 2 reviews the theoretical model underpinning the experimental design, which is described in greater detail in Section 3. Section 4 presents results on pass-through, and Section 5 describes the demand estimation procedure. Section 6 presents the structural estimates of the level of competition among intermediaries and the welfare implications of these findings. Section 7 describes results of the entry experiment. Section 8 concludes.

# 1 Maize Markets in Kenya

Staple commodities represent a major expenditure for consumers across Africa. In Kenya, maize – the country's primary staple commodity – is responsible for over a third of average gross caloric intake. The median household spends 9% of its annual expenditure on maize

[3] The exception is Hildebrandt et al. (2015), which finds that farmers who receive price information earn 5% higher prices for their yams, but this effect disappears by the second year of the study.

[4] In a quasi-experimental variant of this literature, Casaburi, Glennerster and Suri (2013) find that expansion of the road network in Sierra Leone led to price decreases that can be best explained under a search cost framework, and which are inconsistent with either Bertrand competition or Cournot oligopsony.

[5] Observational pass-through rates are much more informative for imported goods or manufactured goods that have a distinct geographic point — and price — of origin

(and the poorest decile spends 14%). On the production side, about half of all Kenyan households grows maize (Argent and Begazo, 2015). The functionality of these staple commodity markets is therefore of significant importance for household welfare.

Figure A.1 displays the maize output market chain in Western Kenya. 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-medium farmers (selling $<$ 5 tons), 16% from large farmers, and 33% from other traders. Traders tend to own a warehouse in a market center and either rent or own a lorry 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 my sample, 64% of sales take place in open-air markets in rural communities. 16% is sold to millers, who grind maize into flour for sale to stores that serve urban consumers. Another 16% is sold to other traders, who sell in other areas of Kenya or eastern Uganda. A small portion of sales – about 2% – is sold to restaurants, schools, and other institutions. Finally, 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.

### 1.1 Entry into Regional Trade

As part of a broad plan of structural adjustment in the 1980s and 1990s, Kenya pulled state-controlled marketing boards out of staple grain markets, lifted trade restrictions on export crops, and allowed prices to be determined by market forces, rather than by state mandate. Today, few legal barriers exist to entering into the maize trade.[6] However, engaging in large-scale, regional wholesale trade still requires significant working capital in order to pay for inventory, storage facilities,[7] and transport vehicles.[8] Further, traders must develop extensive networks of contacts in order to glean information on prices and product availability, as this information is disseminated one-on-one through personal networks of fellow traders rather than through any centralized or open information clearinghouse. It

[6] The few permits that are required are either easy to obtain or are unenforced. The primary license required is the Annual County Business License, which costs about \$100 USD/year and is issued by county officials. Traders report this license is easy to get and most have this license (though most also report that this license is not well-enforced). Other licenses are very poorly enforced, if at all, including a public health license and a transport permit. There are more serious inspections and permits required for cross-border trade. Finally, there is a small \$2 USD "cess" tax charged to traders in the market each day.

[7] Though long-run storage is uncommon among traders, short run facilities are necessary for cleaning, drying, and sorting.

[8] For example, rental of a lorry per day costs \$250 (about 18% of annual GDP per capita), while purchasing a lorry costs \$30,000 (over 21x annual GDP per capita).

is common for traders to enter the business with the support of siblings, spouses, or even former employers who already have experience in the business. Therefore, while entry is close to free legally, those who wish to enter regional trade still face significant barriers.

Table A.1 presents trader demographic details. The average trader has completed some secondary school and is able to answer half of the Ravens matrices (Group B) questions. Only 58% keep written records, which typically include only sale prices and quantities; rarely is cost or accounting data recorded. However, 62% do report reviewing their financial strength monthly. Most traders operate one-man businesses, with only 37% having any employees.

### 1.2 Open Air Markets

This study takes place in the open air markets in which traders sell the majority of their produce. These markets typically occur on a set day each week. The traders present are a mix of those who have their warehouse in that particular market and those who arrive with a truck and sell out of its back for the day. Traders with trucks typically park next to each other in a particular area that they use each week, and warehouses are typically in a row or clustered. Importantly, trader prices, while not posted in any public way, are presumably common knowledge given the close physical proximity of traders. Figure A.2 presents the histogram of the number of traders per market, which varies from 1-10 with a median of 3. Traders commonly work in the same set of markets each week, with 95% of traders reporting working in that market most weeks and only 2% saying that this was their first time in the market (see Table A.1). 77% have worked previously with *all* other traders in the market that day. As a result, 67% say they know the other traders in the market that day "very well," 27% "somewhat well," and only 6% "not very well." When asked directly, only 38% of traders report "discussing a good price" with other traders and only 30% report engaging in an explicit price agreement with other traders; the vast majority claim they are rigorously competing on price. However, 72% of traders work in a market in which at least one trader has reported the existence of a price agreement that day.

Customers in these markets are comprised of two-thirds individual households and one-third rural retailers. The median consumer buys maize only from her local market, though a few retailers purchase from a larger number.[9] I therefore model consumers as being captive to their local market. The median customer buys maize for consumption every week; storage is rare (see Appendix B). The product itself is fairly homogenous (see Appendix C).

[9]This data is drawn from a phone survey with 100 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.

## 2 Theoretical Framework

This study implements three distinct experiments, each of which is designed to identify a specific parameter from a standard model of price setting behavior. Experiment 1 identifies pass-through, while Experiment 2 identifies the curvature of demand. These two parameters are then used to calibrate a structural model of price setting behavior that nests several well-known forms of strategic interaction between traders. With the pass-through rate and demand curvature known, this model enables estimation of a "competitiveness parameter," which reveals the conduct under which traders operate. In the third experiment, the number of traders in the market is experimentally manipulated, and the effect of entry on both conduct and overall competitiveness is estimated. The experimental design is therefore tightly tied to theory. This section reviews that theory.

### 2.1 Model Set-Up

I begin with a standard model of firm profits, in which the profits of a trader in market $d$ on date $t$ can be written as:[10]

$$\pi_{dt} = (P_{dt} - c_{dt})q_{dt} \tag{1}$$

Here, I employ a few simplifying assumptions. First, I assume that maize is a homogenous good and that traders are unable to price discriminate.[11] These assumptions ensure that a single market price prevails and provide a theoretical link between market prices and individual traders' strategic interaction. Second, consistent with Fafchamps, Gabre-Madhin and Minten (2005), I assume marginal costs $c_{dt}$ are constant with respect to quantities. This appears to be a good approximation of the empirical setting, in which the major variable costs are constant with respect to quantity (see Appendix E). Finally, I assume symmetry across traders, specifically with respect to initial marginal cost.[12]

[10] The model employed here is static. While maize is in theory storable, empirically, consumer stockpiling is quite limited (see Appendix B). Related work in the region suggests that credit constraints limit households' ability to arbitrage these price fluctuations (Burke, Bergquist and Miguel, 2016).

[11] Maize sold in these markets is fairly homogenous; there is little variation in quality and credit is rarely used (see Appendix C). Price discrimination appears quite rare, with an intra-cluster correlation of 0.9 between the prices that a trader offers his various customers throughout the day (see Appendix D). This is likely because negotiations between traders and consumers are conducted in public, thereby limiting traders' ability to engage in dramatic price discrimination.

[12] The relative equality of market shares supports this assumption (the variance of market shares – sometimes called the "Asymmetry Index" – is only 5%). Moreover, the feature most crucial to the experimental

Taking the derivative of Equation 1 with respect to the trader's quantity $q_{dt}$ yields the trader's first order condition:

$$P_{dt} = c_{dt} - \theta \frac{\partial P_{dt}}{\partial Q_{dt}} \frac{Q_{dt}}{N_{dt}} \tag{2}$$

where $Q_{dt}$ is the total quantity in the market, $N_{dt}$ is the number of traders in the market, and $\theta \equiv \frac{\partial Q}{\partial q}$ has the following interpretation:[13]

$$\theta = \begin{cases} 0 & \text{when perfectly competitive} \\ 1 & \text{when Cournot competitive} \\ N & \text{when perfectly collusive} \end{cases} \tag{3}$$

Returning to Equation 2, we see that – aside from the shape of demand – prices depend on two features of market structure and trader behavior: the number of traders $N_{dt}$ and how those traders interact according to $\theta$. Following Atkin and Donaldson (2015), I synthesize these two features into a single "competitiveness parameter":

$$\sigma \equiv \frac{N}{\theta} \tag{4}$$

Sensibly, competitiveness in the market goes up with both the number of traders (holding conduct constant) and with more competitive conduct (holding the number of traders constant). I summarize the competitiveness parameter under different models of competition:

$$\sigma \equiv \frac{N}{\theta} = \begin{cases} \infty & \text{when perfectly competitive} \\ N & \text{when Cournot competitive} \\ 1 & \text{when collusive} \end{cases} \tag{5}$$

design is that the *change* in costs is symmetric across traders, which the symmetric experimental manipulation of costs is explicitly designed to ensure.

[13]Under this formulation, $\theta$ is similar to the "conduct parameter" that has fallen out of favor in recent decades for many reasons (Corts, 1999), among which is that it only takes on a clear, well-defined interpretation at the three above values. However, in Section 6, I will primarily use this framework to differentiate between these three well-defined models, by defining the pass-through rate that one should expect to observe under each of these three models and comparing this to the empirical rate. I employ $\theta$ as a convenient formulation for expressing in one nested model the three contrasting forms of competition to be tested.

Because $\sigma$ synthesizes the components of competitiveness, I will work with $\sigma$ in the first portion of this paper, which measures the competitiveness of these markets. However, it is useful to keep the derivation of $\sigma$ in mind when I turn to the effects of entry (i.e., increasing $N$) on conduct $\theta$ and overall competitiveness $\sigma$.

### 2.2 Pass-Through and Demand Curvature

To identify how traders respond to reductions in their marginal costs, taking the derivative of Equation 2 with respect to $c_{dt}$ yields:

$$\rho_{dt} \equiv \frac{\partial P_{dt}}{\partial c_{dt}} = \left\{1 + \frac{1 + E_{dt}}{\sigma_{dt}}\right\}^{-1} \tag{6}$$

where $E_{dt} \equiv \left\{\frac{Q_{dt}}{\frac{\partial P_{dt}}{\partial Q_{dt}}}\right\}\left\{\frac{\partial \frac{\partial P_{dt}}{\partial Q_{dt}}}{\partial Q_{dt}}\right\}$ is the elasticity of the slope of inverse demand. Therefore, the level of pass-through $\rho$ depends on both the competitive structure of markets $\sigma$ and the curvature of demand $E$.

Figure F.1 provides a visual example of this relationship. In the left panel, a cartel determines how much of a $\Delta$ reduction in marginal cost $c$ to pass-on to the price. With moderately curved demand, the cartel will choose to pass on only a fraction of the cost reduction. The right panel presents a market operating under Cournot competition but a more concave demand function. We see that different combinations of competition and demand curvature can yield the same observable pass-through. Therefore, in order to infer the level of competition from pass-through, we must understand the curvature of demand.

### 2.3 Degree of Competition and Welfare Implications

My first experiment estimates pass-through and my second experiment estimates demand curvature. I then use these experimentally estimated parameters to calibrate Equation 6 and back out $\sigma$, the implied degree of competition in these markets.

I can then identify the division of total variable surplus in the market between consumers and intermediaries, as well as deadweight loss, under this market structure. Atkin and Donaldson (2015) solve for the following ratios for consumer surplus (CS), intermediary surplus (IS), and deadweight loss (DWL):

$$\frac{IS}{CS} = \frac{1}{\bar{\rho}} + \frac{1-\sigma}{\sigma} \tag{7}$$

$$\frac{DWL}{IS} = (1-\bar{\rho}) + \bar{\rho}\sigma - \left(\frac{\bar{\rho}\sigma}{(1-\bar{\rho}) + \bar{\rho}\sigma}\right)^{\frac{\bar{\rho}}{1-\bar{\rho}}} (\bar{\rho}\sigma + 1) \tag{8}$$

where $\bar{\rho}$ is the quantity-weighted average pass-through rate.

Intuitively, $\sigma$ summarizes the market structure, while $\bar{\rho}$ (conditional on $\sigma$) summarizes the shape of demand. Together, the two identify the division of welfare in this model. Equations 7 and 8 also allow for counterfactual simulations in which I evaluate the welfare implications of increases in $\sigma$, the degree of competition.

### 2.4 The Effect of Entry on Competition

How will the level of competition $\sigma$ change with entry? It is clear from Equation 4 that as $N$ increases, all else equal, $\sigma$ will increase. However, how entry will affect conduct – that is, the value of $\frac{\partial \theta}{\partial N}$ – is unknown theoretically and must be evaluated empirically. This is what I do in Experiment 3.

## 3 Experimental Design

### 3.1 Sample Selection and Experimental Schedule

The sample of markets in this study is drawn from six counties in Western Kenya. A listing exercise was conducted with the Director of Trade in each county to get a comprehensive list of all markets in the county. Markets without maize traders and urban markets in town centers were then excluded. See Appendix G for additional details on the sample selection procedure.

The two market-level experiments (pass-though and entry) were each run for four weeks in a row. This time spans about a quarter of the full selling season in the region (March to July). This duration of treatment was selected to represent a long-run cost or entry shock. It was also selected to match the frequency at which prices regularly vary to minimize concerns about sticky prices (see Figure 1, which displays the relative size of the subsidy compared to weekly fluctuations in market prices). Because piloting revealed that market and week fixed effects were important (cutting standard errors almost in half), the experiment was designed to provide each market each treatment status (pass-through treatment, entry treatment, and control) in a random order to allow for the inclusion of these fixed effects. Figure G.1 shows the six possible orders.[14] Each four-week block was broken by a one-week break during which the demand experiment was run in a subset of markets.

### 3.2 Experiment 1: Pass-Through

In treatment market-days for the pass-through experiment, all traders in the market were offered a subsidy per kg sold. Enumerators arrived at the market at 7:30am (prior to the market start) and immediately made the offer to every trader present. Any traders who arrived later were also presented with the offer immediately upon arrival. Enumerators stayed in the market until 5pm (after the market conclusion). Maize sold during the enumerators' presence in the market was eligible for the subsidy.[15] 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."

In the first week, traders were informed that the offer would be available for four weeks. An identical script was read in each subsequent week to remind returning traders of the availability of the subsidy and to make the offer to any new traders who were absent in the previous week. All traders in the market therefore received an identical reduction in their marginal costs, a crucial feature to map the experiment to pass-through theory.

The 60 markets in the sample were divided into two groups: 45 markets received a "low" subsidy level of 200kg/90kg bag when they were in the pass-through market treatment and 15 markets received a "high" subsidy level of 400kg/90kg bag (sales of partial bags were eligible at the same prorated amount). Note that "low" and "high" are merely relative titles: both represent large and meaningful changes to traders' costs. The "low" subsidy rate represents 7.5% of the average price, while the "high" subsidy represents 15% of the average price. Payments were made via mobile money twice a day.

[14]This randomization was first blocked by the day of the week of the market (done primarily for logistical ease as the pass-through and entry treatment required additional management time to facilitate payments, and equal distribution of treatment across days of the week ensured an even flow of management duties) and then stratified by the number of traders typically in the market, as identified in the market census. See Appendix G for further details on this census.

[15]Only maize sold in cash was eligible for the subsidy due to concerns about the ability of enumerators to verify the authenticity of credit sales. However, over 95% of sales are conducted in cash, so this restriction was often irrelevant. The subsidy was capped at the first 75 90kg bags sold to limit budget exposure, but this cap was binding for only 1.5% of traders.

Enumerators monitored the sales of each trader throughout the day, recording the price and other details of each transaction as will be described below in the data section. The data collection process was identical in treatment and control markets.

### 3.3 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.[16] 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. 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.

Traders' consent was acquired at the beginning of each day. The trader was therefore aware that his customers would (potentially) receive price reductions. While this may have changed the baseline price charged by the trader (e.g., the trader may have raised his overall price to collect some of the anticipated discount), the trader did not know at the time of price negation with any one consumer the amount of the discount that would be offered nor was the trader permitted to adjust the price following the announcement of the realized discount amount. Therefore, any variation in price driven by the discount is random.

Discounts were given per kg purchased (so as to lower the price/kg). Ten levels of discounts were offered, calibrated to span the range of price reductions one would have observed if 0-100% of the cost-reduction subsidy had been passed-through in the pass-through experiment. Per 90kg bag, they were: {0, 25, 50, 100, 150, 200, 250, 300, 350, 400} Ksh.

### 3.4 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. Three traders were given the offer for each market. This was designed (1) to increase the probability that at least one trader took

[16] The sample is therefore drawn from consumers who were already planning on purchasing maize that day. This was done out of necessity, in order to identify a pool of "customers" in which to randomize the discount amount. However, it does mean that the sample does not include customers who may have been induced on the extensive margin into market participation at these lower, discounted prices. The assumption therefore in the demand analysis is that these customers would have exhibited the same curvature of demand as the customers observed in the sample.

up the offer and (2) to measure traders' willingness to enter, as the amount of each offer was randomized. Offers were given for four weeks in a row to generate somewhat long-run entry.

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, I 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, I then randomly selected the three traders who would receive the entry offers.[17]

Once the set of offers was established, each of the three selected traders for each market was randomized into a "low" offer of 5,000Ksh ($50 USD), a "medium" offer of 10,000Ksh ($100 USD), and a "high" offer of 15,000Ksh ($150 USD). The trader was eligible to receive this amount each time he visited the entry market on any of four offer days.[18] Payout was contingent on a few factors, of which traders were made aware during the offer call. They were that the trader must: (1) come to the specified market on the specified date; (2) arrive with a truck and at least 15 bags; (3) stay for at least one hour and show intention to attempt sales. Payment was made via M-Pesa immediately after these conditions were met.

Traders were informed of the offer via phone call one week prior to the first market-day for which they were eligible. During this call, a short survey was conducted to gather additional information about the potential entrant, including whether he had contacts in the market, his expected profits for the day should he take up and not take up the offer respectively, and his ethnicity. Following each offer week, four short follow-up phone surveys were conducted, in which information was collected about the trader's activities on the day of the offer regardless of whether or not they accepted the offer.

[17]Because I did not want to overwhelm a single trader with too many offers, I only offered each trader one offer per 4-week block. Because this has cascading effects for the set of eligible traders for each market, I 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.

[18]Traders were encouraged to attend all four days to receive four payouts of the above amounts. Offers for each day were independent because making payouts contingent on perfect attendance could have potentially discouraged overall take-up.

### 3.5 Data

Data was collected in an identical way in all markets and in all periods (pass-through treatment, entry treatment, and control). Depending on the activity level of each market, enumerators were assigned to monitor 1–4 traders.[19] Those surveys captured transaction-level price, quantity, payment method (cash or credit), and customer type (individual household consumer or retailer), all observed in real-time by the enumerator. Data on the value of any non-traditional reductions in price were also collected; these included: flat reductions in the total cost of the purchase (rather than in the per-unit price); "top-ups," quantities of maize added to the total purchase "for free"; and "after-bag services," such as free sacks, transport, or other services given to customers bundled with their transactions. The value of these additional services is incorporated into the price per kg.[20] Maize quality data was also collected for each trader (see Appendix C for greater detail). In addition, traders were asked about their experience with other traders in the market that day: how often they had worked with others before, how well they knew others, whether they had "discussed a good price" at which to sell, and whether they had "agreed on a price" at which to sell.[21] Finally, the first time a trader was met in the sample, additional information was captured on the trader's fixed characteristics, including ethnicity, location of home market, highest level of education achieved, and a battery of business management and record keeping questions drawn from McKenzie and Woodruff (2015). A Raven's test was also administered.

My primary outcome of interest – price – is defined as the quantity-weighted average of transaction level prices that the trader sold that day:

$$P_{idw} = \frac{\sum_{t=1}^{T} p_{idwt} q_{idwt}}{\sum_{t=1}^{T} q_{idwt}} \tag{9}$$

where $p_{idwt}$ is the price of transaction $t$ for trader $i$ in market $d$ in week $w$ and $q_{idwt}$ the quantity (though note that because the ICC of price within a trader in a given market-day is

[19] Busier markets with more quickly moving sales were allocated additional enumerators to ensure that all transactions could be recorded with accuracy.

[20] These non-traditional reductions in price were uncommon, but they do add 1–2 percentage points to my measure of pass-through, so there is some indication that traders can use these less-traditional methods of price reductions to pass-through some of the cost reduction. It is possible that this is a more discrete method of deviating from price agreements maintained with fellow traders.

[21] Due to their sensitivity, these questions were asked at mid-day, after the enumerator had established good rapport with the respondent. For any traders who left the market before that time, enumerators attempted to ask these questions before the trader left, but these efforts occasionally failed due to short notice. As a result, there is higher attrition among this section of the survey.

high (0.9), in practice there is little variation in the prices entering into this average). This measure of price per trader $P_{idw}$ forms the basis for the primary analyses of the pass-through and entry experiments. All estimates are weighted by the inverse of the number of traders in the market so as to give equal weight to each market in the final analysis. All standard errors are clustered at the level of the market-block, the unit of randomization.

# 4 Pass-Through

To measure pass-through, I estimate:

$$P_{idw} = \alpha + \beta CR_{dw} + \boldsymbol{\gamma_w} + \boldsymbol{\zeta_d} + \epsilon_{idw} \tag{10}$$

where $P_{idw}$ is the average price per kg charged by trader $i$ in market $d$ in week $w$, $CR_{dw}$ is the level of cost reduction per kg offered in market $d$ on week $w$ (i.e., CR is the *negative* value of the marginal subsidy in pass-through treatment markets and zero elsewhere), $\boldsymbol{\gamma_w}$ and $\boldsymbol{\zeta_d}$ are week and market fixed effects, respectively, included to improve precision. The sample includes traders in market-days in which the market was in either the pass-through treatment or control period – market days assigned to the entry treatment are omitted. Under this specification, the coefficient of interest is $\beta$, which yields the pass-through rate, or $\frac{\partial P}{\partial c}$.

To measure heterogeneity in the pass-through rate by the level of the cost-reduction, I estimate

$$P_{idw} = \alpha + \beta_1 CR_{dw} * Low_{dw} + \beta_2 CR_{dw} * High_{dw} + \boldsymbol{\gamma_w} + \boldsymbol{\zeta_d} + \epsilon_{idw} \tag{11}$$

in which $Low_{dw}$ ($High_{dw}$) is a dummy indicating whether the market was in a low (high) subsidy market. This allows for non-linearities in the effect of the subsidy per kg. For other measures of heterogeneity, I run specifications similar to Equation 11, conditioning on the desired dimension of heterogeneity.

Table 1 presents the main results of the pass-through experiment. In Column 1, I see that pass-through is 22.4%, significantly different from zero at the 1% level and measured with a high degree of precision. Column 2 presents pass-through rates for low and high cost reduction treatments separately. The pass-through rates for each group are almost identical. This constant empirical pass-through rate will provide important empirical justification for the functional form assumptions in the following section on demand estimation.

I explore heterogeneity by the number of traders in the market.[22] Figure 2 presents these results, which show little evidence of meaningful heterogeneity. Estimates of pass-through rates are fairly tightly centered around the overall estimate of 22% and no clear pattern is seen with the number of traders. To gain statistical power, the bottom two measures show the sample pooled into below- and above-median number of traders; again, point estimates are not statistically significantly different and are in fact remarkably close in magnitude.

I further explore other sensible dimensions of heterogeneity by a few other measures in Figure 3.[23] First, I measure whether pass-through is different for markets on and off tarmac roads, which serve as a proxy for market geographic isolation. I find no evidence of heterogeneity by this measure. Next, I explore whether a higher intensity of explicit collusion predicts lower pass-through rates, measured by looking at the number of market-days within a market where traders have explicitly admitted to collusion.[24] The point estimates suggest that pass-through is sensibly smaller for markets above the median in this measure, but these differences are neither statistically significant nor large in magnitude. In summary, the lack of clear heterogeneity and relatively consistent point estimates suggests that pass-through is fairly constant across markets.

## 5 Demand Estimation

As described in Section 2, in order to draw inference about the level of competition from the observed pass-through, one must first understand the curvature of demand. To do so, I use the demand experiment to estimate a general Bulow-Pfleiderer class of demand functions:

(12)

$$Q_{dt}(P_{dt}) = \begin{cases} \left(\frac{a-P_{dt}}{b}\right)^{\frac{1}{\delta}} & if\ \ (P_{dt} \leq a, b > 0\ \ and\ \ \delta > 0)\ \ or\ \ (P_{dt} \geq a, b < 0\ \ and\ \ \delta < 0) \\ 0 & if\ \ P_{dt} > a, b > 0\ \ and\ \ \delta > 0 \\ \infty & if\ \ P_{dt} \leq a, b < 0\ \ and\ \ \delta < 0 \end{cases}$$

$where\ a \geq 0$

[22] This is the main source of heterogeneity pre-specified in a design registry submitted prior to the beginning on the experiment. The number of traders is defined as the average number of traders observed in the market over the course of the experiment. In order to remove any increases in the number of traders driven by the entry experiment, this figure uses the average of the predicted number of traders each week, based on market and week fixed effects.

[23] These were not included in the design registry.

[24] I construct, for each market, a count of the number of market-days in which at least one trader admitted to discussing (agreeing on) prices with other traders. I then divide the sample into markets above and below the median of this measure.

I choose this particular class of demand functions for its flexibility, tractability, and empirical foundation. First, this demand structure is flexible because it nests many of the functional forms common to the development and trade literature, including linear demand, quadratic demand, and isoelastic demand, rather than assuming a particular functional form.

Second, this class of demand functions is tractable, producing a constant elasticity of the slope of inverse demand with respect to quantity ($E$) (Bulow and Pfleiderer, 1983). Recall that in the model described in section 2, the pass-through rate is determined by the competitiveness parameter $\sigma$ and the slope of inverse demand $E$. While in theory $E$ can vary with $q$, this term is a second order term for which it is already difficult to get precision on a single estimate using the full pooled data (as I will show below); attempting to further estimate $E$ at different levels of $q$ would be even more challenging. Under the Bulow-Pfleiderer class of demand functions, $E$ is constant with respect to $q$. To see this, note that the inverse demand function is:

$$P_{dt} = a - bQ_{dt}^{\delta} \tag{13}$$

In this case, the elasticity of the slope of inverse demand, $E_{dt} \equiv \left\{ \frac{Q_{dt}}{\frac{\partial P_{dt}}{\partial Q_{dt}}} \right\} \left\{ \frac{\partial \frac{\partial P_{dt}}{\partial Q_{dt}}}{\partial Q_{dt}} \right\}$ reduces to $\delta - 1$. Therefore, Equation 6 simplifies to:

$$\rho_{dt} \equiv \frac{\partial P_{dt}}{\partial c_{dt}} = \left( \frac{\sigma}{\sigma + \delta} \right) \tag{14}$$

Third, this class of demand functions has a strong empirical foundation. The experimental design includes variation intentionally designed to test this empirical fit. As shown by Equation 14, because $E$ is constant across $q$, this class of demand functions predicts a constant pass-through rate for a given $\sigma$, independent of the size of the cost shock (were $E$ not constant in $q$, cost shocks of different sizes – by driving different levels of optimal quantity sold – would induce differential changes in $E$, which would in turn produce different pass-through rates). By offering two different levels of the cost shock, I am able to test for this prediction of constant pass-through. Because markets are randomized into receiving the low vs. high subsidy rate, one can assume these two sets of markets have, on average, identical levels of competitiveness ($\sigma$). Therefore, the only difference in these two sets of markets, on average, should be the level of the cost shock. Under the Bulow-Pfleiderer class of demand functions, we should therefore expect to see identical pass-through rates for these two markets. This is exactly what we see in Column 2 of Table 1, which suggests remarkably similar pass-through rates for the two levels of cost reduction. This lends empirical support to this choice of demand class.

### 5.1 Estimation and Results

I utilize the randomized reduction in the price paid by consumers from the demand experiment as an instrument for price. The analysis is run with 1,206 observations. I estimate the vector of parameters $\boldsymbol{\Theta} = (a, b, \delta)'$ in Equation 12 using generalized methods of moments with a vector of sample moments given by $\boldsymbol{m(\Theta)} = \boldsymbol{Z'\xi(\Theta)}$. Here, $\boldsymbol{Z}$ is a matrix of instruments formed by the stacked row vectors $\boldsymbol{Z_i} \equiv (1, d_i, d_i^2)$, with $d_i$ defined as the value of the discount amount randomly offered to customer $i$ (recall $d$ is one of the ten possible discount values). The vector $\boldsymbol{\xi}$ is the stacked residuals from a logged transformation of Equation 12 such that $\xi_i = logQ_i - \frac{1}{\delta}log(a - P_i) + \alpha$, where $\alpha \equiv \frac{1}{\delta}log(b)$. Thus, the parameter estimates are given by the GMM objective function:

$$\boldsymbol{\Theta}^* = \underset{\boldsymbol{\Theta}}{\operatorname{argmin}}\, \boldsymbol{m(\Theta)'Wm(\Theta)}, \tag{15}$$

which is estimated in two steps: the first step with the weighting matrix $\boldsymbol{W} = (n^{-1}\boldsymbol{Z'Z})^{-1}$ and the second step, in which the weighting matrix is replaced with the estimated optimal weighting matrix $\boldsymbol{W} = (\frac{1}{N}\boldsymbol{g'g})^{-1}$, where $\boldsymbol{g}$ is a matrix formed by the stacked row vectors of $\boldsymbol{g_i = Z_i\xi_i(\Theta_1)}$.

Because there are two sets of possible constraints on the parameters in order to see positive, finite demand, I estimate the model under each set of constraints separately. I find that the minimand is smaller under the first set of constraints and so continue under this set of constraints.[25] Moreover, note that the second set of constraints, in which $\delta < 0$, would suggest pass-through rates of *greater* than 100% under imperfect competition, which is inconsistent with what is observed in practice (though it is important to emphasize that the demand estimation is in no way constrained by the results of the pass-through experiment).

Estimates are initialized at 500 randomly selected starting values, to ensure the min-

[25] The minimum of the objective function achieved under the first set of constraints is $5.7x10^{-4}$, while under the second it is $2.5x10^{-3}$. Note also that under the second set of constraints, the point estimate on $\delta$, the parameter of interest, is very close to the bound of 0 ($\hat{\delta} = -0.0964$, with a large standard error of $1.4x10^{4}$).

imization procedure does not obtain parameters for a local minimum[26]. I then generate bootstrapped confidence intervals by estimating these parameters on 1,000 random draws (with replacement) of the data.

Results are presented in Table 2, which show the point estimate and 95% confidence interval. Note that the confidence interval on $\delta$ is wide. For example, I cannot rule out linear demand ($\delta = 1$), nor can I rule out very curved inverse demand ($\delta = 5.89$). This is because $\delta$, which represents the elasticity of the slope of inverse demand (plus one), is a higher order object which I am underpowered to measure with great precision, even with over 1,200 observations from the demand experiment. However, we will see in the next section that even this limited precision is sufficient for our purposes. From the point estimate on $\delta$, I can predict the level of pass-through that one should expect under various models of competition; I will find the prediction of one model to line up very closely with what is observed empirically. Moreover, even at the bounds of my estimate on $\delta$, I can still reject that what I see empirically is consistent with other common models of competition.

# 6 Degree of Competition and Welfare Implications

First, I demonstrate that the observed pass-through is very close to the collusive model prediction evaluated at a demand curvature given by the parameter point estimates. Given the point estimate on $\delta$ of 2.8, I use Equation 14 to estimate the average pass-through rate one should expect to observe in the experiment under various models of competition. If markets are perfectly competitive ($\sigma = \infty$), we should observe 100% pass-through. If markets are Cournot competitive ($\sigma = N$), we should observe pass-through rates that vary with the number of traders: $\rho = \frac{N}{N+2.8}$.[27] Given the distribution of number of traders in each market in my sample, the expected pass-through rate if markets are Cournot competitive is 55%. Finally, if markets are collusive ($\sigma = 1$), we should expect to observe 26% pass-through.

Figure 4 displays the bootstrapped distribution of $\rho$.[28] I see that the mass of the dis-

[26] These values are drawn from a uniform distribution spanning the range of feasible parameter estimates. For example, for $\delta$, the primary parameter of interest, the range of start values ranges from $e^{-10}$ to $e^{10}$. These (more than) span the values of $\delta$ that represent linear demand ($\delta = 1$) to extremely curved demand. Most importantly, the range of possible $\delta$'s span those that would reconcile the observed pass-through rate of 22% with the full set of models considered here, from perfect collusion to near-perfect competition and therefore allow differentiation between the set of market structure models considered here. I again emphasize that the demand estimation is in no way constrained to match any moments from the pass-through experiment.

[27] This would predict that pass-through would be increasing in the number of traders. Note that we already saw in Section 4 that pass-through did not vary with the number of traders in way that is consistent with this predicted pattern.

[28] The distribution was constructed using 1,000 block bootstrapped samples where blocks are defined by market by 4-week treatment-blocks. There are 180 such clusters from 60 markets.

tribution of $\rho$ is concentrated near the predicted pass-through of 26% under collusion. The dotted lines, which identify the 90% confidence interval, clearly reject a $\rho$ consistent with that predicted under a model of Cournot competition or perfect competition.

This exercise does not take into account the fact that $\delta$ is estimated imprecisely. To account for this imprecision, I generate a bootstrapped distribution of $\sigma$ by calibrating Equation 14 with 1,000 bootstrapped estimates of $\rho$ and $\delta$. Figure 5 presents this distribution, overlaid with the benchmark values of $\sigma$ under each model of competition.[29] I plot in red the value of $\sigma$ predicted by the point estimates on $\rho$ and $\delta$. The point estimate of $\hat{\sigma}$ is 0.81, which is quite close to – and statistically indistinguishable from – the model benchmark of $\sigma = 1$ under perfect collusion. Moreover, while the collusive market benchmark of $\sigma = 1$ lies squarely in the middle of the 90% confidence interval, the levels of $\sigma$ predicted by a Cournot model and perfectly competitive model lie outside these bounds. I am therefore able to reject them with 90% confidence.[30]

The observed pass-through rate is therefore consistent with an underlying market structure in which traders exert a high degree of market power. The following section describes the welfare implications of this lack of competition.

### 6.1 Welfare Implications

What does this imply for the division of surplus between consumers and intermediaries? I use Equations 7 and 8 to solve for the ratios for consumer surplus (CS), intermediary surplus (IS), and deadweight loss (DWL).[31] Table 3 shows the results. At $\sigma = 1$ (the closest model-consistent value to the estimated $\sigma$ of 0.81) and a $\rho$ of 0.26 (the $\rho$ which would be consistent with this $\sigma$), I estimate that only 17.8% of the total variable surplus generated by the maize market accrues to consumers, while intermediaries reap 67.6%. Another 14.6% is lost to DWL. Even at the upper edge of the confidence interval around $\sigma$ (and the corresponding $\rho$), consumers are at most receiving 25.5% of the surplus. Therefore, I see that intermediaries accumulate much of the gains from these transactions.

[29] Note that for Cournot $\sigma = N$ and therefore varies across markets, I show the average $\sigma$ we should expect to see if all markets are behaving in a Cournot competitive manner, given the distribution of the market sizes observed in my sample.

[30] I am able to reject a Cournot competitive model because the confidence interval around $\delta$, however large, does exclude the extreme curvature necessary to justify such low pass-through under a Cournot model. To achieve a predicted $\rho$ of 22% under a Cournot model, we would have required a $\delta$ of about 12. For the perfectly competitive model, the predicted $\sigma$ of $\infty$ lies all the way to the right outside the range of the figure and is clearly rejected.

[31] Under the assumption of Bulow-Pfleiderer demand, which implies a constant pass-through rate, $\bar{\rho}$ collapses to $\rho$.

I can also conduct welfare counterfactuals by calibrating Equations 7 and 8 with the value of $\sigma$ that corresponds to counterfactual forms of market conduct and the $\rho$ that would be realized at each of these values of $\sigma$.[32] Table 4 presents the results of this exercise. I find that if markets were Cournot competitive, consumers would reap 49% of total variable surplus, and if markets were perfectly competitive, they would receive 100%.[33] Part of this gain in consumer surplus is a transfer from intermediaries to consumers, but this may be in keeping with the preferences of a policymaker who places greater weight on the welfare of poor rural consumers than on intermediaries. The reduction in deadweight loss is an unambiguous gain. Figure 6 presents the same results in a more continuous form.

Increasing competition among intermediaries would therefore yield large welfare gains for consumers, could such a goal be achieved. It is this goal that I address in the next section.

# 7 Generating Entry

Given that markets look fairly collusive, one natural policy response is to encourage greater entry. There are several policies that could potentially encourage entry, such as offering lines of credit to potential new traders to rent long-haul trucks, disseminating information about good markets more broadly, etc. However, it is unknown how much entry will enhance competition and improve consumer welfare. This is what I measure in the third experiment, in which I randomly incentivize traders to enter new markets.

## 7.1 The Cost of Entry

The cost of entry appears to be high in this setting. Because the offer amount is randomized, I can use traders' willingness to accept the offer as a measure of willingness to enter new markets. Table 5 presents take-up at each subsidy level (take-up defined as ever accepting any of the four market-day offers). Sensibly, I see that take-up increases in the size of the subsidy: take-up is 12% for the low offer, 28% for the medium offer, and 42% for the high offer. However, these rates are low compared to the percentage of traders who report that it would be profitable to take-up the subsidy given their offer size: 77% at the

[32] As estimated by Equation 14, using the counterfactual $\sigma$ and estimates of $\hat{\delta}$.

[33] This, along with the other welfare results, relies on the assumption of constant marginal costs. If marginal costs were increasing in quantity, intermediates would reap some positive percentage of the surplus in a competitive environment (however, as documented in Appendix E, the empirical evidence is consistent with constant marginal costs). Relatedly, if traders were to price at average cost under perfect competition, intermediates would also earn a positive percentage of the surplus equal in absolute magnitude to their fixed costs.

low offer, 80% at the medium offer, and 89% at the high offer.[34] This discrepancy – the fact that low take-up is observed despite the fact that traders themselves report that take-up would for the most part be monetarily profitable – hints at the existence of barriers or forces discouraging traders from entering otherwise profitable markets.[35]

In order to shed light on these potential barriers to entry, as well as to understand general variation in willingness-to-enter, I explore heterogeneity in offer take-up by a few key variables pre-specified in the design registry. While these results are merely correlational, and therefore cannot be interpreted through a strictly causal lens, they do point to some potential barriers to entry. To explore this heterogeneity, I estimate the following regression specification on the pool of 180 potential entrants:

$$\begin{aligned} T_{id} &= \alpha + \beta X_{id} + \epsilon_{id} \quad (i) \\ T_{id} &= \alpha + \beta X_{id} + \boldsymbol{\zeta_d} + \epsilon_{id} \quad (ii) \end{aligned} \tag{16}$$

in which $T_{id}$ is a indicator representing whether trader $i$ ever took up an offer to enter his assigned market $d$. $X_{id}$ is the variable by which I explore heterogeneity. In specification (ii), I control for market fixed effects ($\boldsymbol{\zeta_d}$), such that I only look at differential take-up of the entry offer *within* the same market. I do this to remove some of the endogeneity that might influence the composition of the pool of potential entrants. Because there were a few traders who were given multiple offers (though never for the same four-week block), I cluster standard errors by trader in both regressions.

Figure 7 displays the results. As presented earlier, a larger subsidy increases take-up. Longer distances to travel are also sensibly correlated with lower take-up; when comparing distance's effect on take-up with the offer amount, an additional 50km in distance is roughly equivalent to a drop of $46 USD in the offer amount.[36] Having contacts in the entry market is correlated with higher take-up (albeit not quite significantly). The point estimate suggests

[34] In a survey conducted during the offer phone call, traders reported the revenues and costs they would expect to incur if they did not take up the offer and instead followed their typical schedule and similarly, the revenues and costs they would expect to incur if they did take up the offer. When the profits expected under take-up plus the offer amount exceed the profits expected under no take-up, I code this as reporting that take-up would be profitable.

[35] Low take-up could also be due to trader mistrust of the offer. However, Innovations for Poverty Action (IPA), the implementing partner, had been conducting surveys with traders in the region for almost three years at the time of the experiment and therefore was well-known by many of these traders. As a result, when asked, less than 5% of traders who did not take up the offer cite trust issues as the explanation.

[36] The magnitude and precision of the distance effect drop when including market fixed effects; this is likely because comparing variation in distance to the same market removes much of the total variation in distance.

that the value of having contacts is equivalent to an increase in the offer amount of $36. Being a large firm (above median profits) is also correlated with higher take-up. The effect is substantial: having above median profits is equivalent to offering an additional $52. These results on contacts and firm size are consistent with the existence of barriers to entry in the form of requiring business networks and access to working capital to enter new markets.

Interestingly, ethnic similarity between potential entrants and incumbents does not appear to have any correlation with the entrant's willingness to enter. This is perhaps surprising, given recent work from the region documenting the important role ethnic divisions can play in discouraging productivity among workers (Hjort, 2014) and integration across markets (Robinson, 2016). However, it is consistent with economic lab games from Kenya that fail to find evidence of co-ethnic bias, instead suggesting that observed ethnic divisions may be caused by mechanisms other than simple ethnic preferences (Berge et al., 2015).

Because the offer was made to three different traders per market, this offer generates a strong instrument for entry (despite the low take-up per trader). 53% of all markets had at least one day (out of four) with entry. 38% of all market-days had entry. And 26% of all market days had more than one entrant. In total, an average entry market had an additional 0.6 traders present, an increase of 13% over the mean market size and 20% over the median.[37]

### 7.2 The Effect of Entry on Price

I turn now to the effect of entry on prices. To measure the reduced form effect of the offer, I estimate:

$$logP_{idw} = \alpha + \beta EOM_{dw} + \boldsymbol{\gamma_w} + \boldsymbol{\zeta_d} + \epsilon_{idw} \tag{17}$$

where $logP_{idw}$ is the log of the average price per kg charged by trader $i$ in market $d$ in week $w$, $EOM_{dw}$ ("Entry Offer Market") is a dummy for whether market $d$ is in an entry market in week $w$, and $\boldsymbol{\gamma_w}$ and $\boldsymbol{\zeta_d}$ are week and market fixed effects respectively. Standard errors are clustered at level of market x four-week block, the level of randomization. Observations are weighted by the inverse of the number of traders in each market to give each market equal weight. The sample includes traders in market-days corresponding to either the entry treatment or control period (that is, pass-through treatment periods are omitted). Under

[37]Appendix H documents how entrants compare to incumbents in the same market. I do not see any statistically significant differences in terms of quantity sold or price at which sold between the entrants and incumbents, though point estimates suggest that entrants may sell slightly less and at a slightly lower price.

this specification, the coefficient of interest is $\beta$, which yields the percent reduction in price observed in the entry offer market.

I also run a similar specification to determine the effect of entry on prices:

$$logP_{idw} = \alpha + \beta \hat{N}_{dw} + \boldsymbol{\gamma_w} + \boldsymbol{\zeta_d} + \epsilon_{idw} \tag{18}$$

in which $\hat{N}_{dw}$ represents the number of traders in the market that day, for which I instrument with the $EOM_{dt}$ dummy. Table 6 presents these results. Despite a strong first-stage effect on the number of traders (Column 1), reduced form effects are small and not quite significant, with only a 0.6% drop in prices (Column 2). Column 3 presents the result of using treatment status as an instrument for the number of traders. I see that the entry of one trader reduces pries by 1% (p-value of 0.101).

Figure 7 presents heterogeneity in entry effects along different dimensions of market characteristics, which should be interpreted with caution.[38] Panel A presents differences in take-up rates by markets with above vs. below median number of traders, markets on vs. off tarmac (paved) roads, markets with above vs. below the median number of reports of price discussions and markets with above vs. below the median number of reports of price agreements. No clear differences in take-up are seen across these groups, with the exception of markets with a greater number of traders, which do have statistically significantly higher take-up. Panel B presents IV effects on price broken down by the same categories. No definitive patterns of heterogeneity emerge based on the number of traders in the markets (see Table I.1 for further breakdown of these effects) or whether the market is on tarmac road. However, it does appear that what small decreases I do observe in price are concentrated in markets in which fewer traders report discussing or agreeing on price.

### 7.3 The Effect of Entry on Competition

Given that I observe a reduced form price decrease of 0.6%, what does this tell us about how the underlying competitive environment ($\sigma$) has changed? Recall that $\sigma = \frac{N}{\theta}$. The effect of the experiment on $N$ is directly measurable – this is the first stage effect of 0.6 – but what is unknown is how entry will affect the conduct between traders $\theta$.

[38]Recall that the pool of potential entrants differs by market, necessitated by the requirement that one has never worked in that market before. It is likely that variation in market characteristics are correlated with variation in characteristics of the entrant, and it is therefore difficult to separate what differences in observed effects are due to variation in market characteristics versus entrant characteristics. That said, from the policymaker's perspective, separating these two may not be crucial if the two are correlated in practice.

To explore this, I consider three possible scenarios for how one would expect $\theta$ – and therefore $\sigma$ and ultimately prices – to change with this degree of entry:[39]

1. No change in $\theta$: **conduct unchanged**. The effect on competitiveness $\sigma$ just the mechanical effect of raising $N$ by $\Delta N$. In this case, $\theta$ remains equal to $N_0$, and the new $\sigma = \frac{N_0+\Delta N}{N_0}$

2. Decrease in $\theta$: the **entrant competes** with the incumbents. In this case, incumbents continue to act as a block, such that $\theta_I = N_0$, but entrants act as a competing firm, such that $\theta_E = 1$. In this case, the average $\theta$ in the market becomes $\frac{N_0^2+\Delta N}{N_0+\Delta N}$ and the average $\sigma = \frac{(N_0+\Delta N)^2}{N_0^2+\Delta N}$

3. Increase in $\theta$: the entrant simply **joins the cartel**. In this case, $\theta$ increases by $\Delta N$ to offset the increase in $N$, leaving market competitiveness $\sigma$ unchanged. $\theta = N_0 + \Delta N$ and $\sigma = 1$

What price effects should one expect at these various levels of $\sigma$? Returning to theory, recall that that trader's first order condition for prices is:

$$P_{dt} = c_{dt} - \theta \frac{\partial P_{dt}}{\partial Q_{dt}} \frac{Q_{dt}}{N_{dt}} \tag{19}$$

Assuming $\mathrm{E}[c_T - c_C] = 0$ (which is true by construction of the RCT, at least for incumbents, and empirically true for entrants as well):[40]

$$E[P_T - P_C] = E\Bigg[ - \left(\frac{Q_T}{\sigma_T}\right)\left(\frac{\partial P_T}{\partial Q_T}\right) + \left(\frac{Q_C}{\sigma_C}\right)\left(\frac{\partial P_C}{\partial Q_C}\right)\Bigg] \tag{20}$$

With Bulow-Pfleiderer demand $\big(Q\big)\big(\frac{\partial P}{\partial Q}\big) = -b\delta Q^{\delta}$. Substituting in $Q^{\delta} = \frac{a-P}{b}$, this simplifies to $\big(Q\big)\big(\frac{\partial P}{\partial Q}\big) = -\delta(a-P)$, which yields:

$$E[P_T - P_C] = E\Bigg[\left(\frac{\delta(a-P_T)}{\sigma_T}\right) - \left(\frac{\delta(a-P_C)}{\sigma_C}\right)\Bigg] \tag{21}$$

[39] An obvious fourth scenario is one in which entry further breaks up collusion among incumbents, which could occur if existing collusive agreements among incumbents become less tenable in the presence of entry. This would produce an even greater price decrease than that expected in scenario 2. Because I do not observe price changes even consistent with scenario 2, I do not consider this scenario in great detail here.

[40] I attempt to measure the major costs faced by traders, such as inventory purchase price, transport costs, etc., and do not see a statistical difference between those of the entrants and incumbents.

The observed price change therefore reflects underlying changes in the competitiveness parameter $\sigma$ (as well as any shifts along the demand curve as prices move). I can therefore look at this problem in two ways. First, taking the point estimates on $\delta$ and $a$ seriously, I can evaluate how much one would expect prices to move with entry for the various potential expected $\sigma_T$. Table 7 presents these simulations. The top panel presents the simulated effect of entry by one trader (the instrumental variable effect) for each market size (as determined by the baseline number of traders). The middle panel presents the reduced form effects to be expected given the first stage effect ($\Delta N$) observed for each market size. This is my preferred benchmark for the expected reduced form effects because the variation in the first-stage effects enters non-linearly into Equation 21.[41]

Using the predicted price change for each market-size, I calculate the average price change one should expect to see under each scenario of entrant behavior given the distribution of market sizes in the sample:

$$\overline{\Delta P^{EB}} = \frac{\sum_{s=1}^{10} \Delta P_S^{EB} N_s}{\sum_{s=1}^{10} N_S} \tag{22}$$

in which $\Delta P_S^{EB}$ is the change in price expected in a market of size $S$ if entrants act according to entrant behavior $EB \in \{$*conduct unchanged, entrant competes, entrant colludes*$\}$. Therefore, $\overline{\Delta P^{EB}}$ identifies the average reduced form price effect one should expect to observe under each model of entrant behavior. I predict a $\overline{\Delta P^{EB}}$ of -4% if conduct is unchanged, -7% if the entrant competes, and 0% if the entrant colludes with incumbents.

The reduced form effect observed of -0.6% is clearly closest to the scenario in which the entrant colludes. Figure 9 presents a graphical version of this intuition. The bootstrapped distribution of the reduced form effect on log price is shown.[42] Overlaid is the reduced form effects on log price that one would have expected if the entrant competes, if conduct remains

[41] However, there is a trade-off here, as there is likely more noise in the first-stage take-up estimates when broken down by market size and this noise also enters non-linearly into Equation 21. As a robustness check, I also run this analysis using the average first stage effect of 0.58, pooled across markets of all sizes, which reduces this noise but also reduces real variation in the first stage ($\Delta N$). Simulation results for this alternative specification are presented in the bottom panel of Table 7. I find that the average predicted price effects under this specification are actually larger than those predicted using per-market-size variation in the first stage (because the first stage is greater in large markets and the IV effect of entry is smaller in large markets). Because I am able to reject the smaller levels of price changes predicted by the main specification, I am also able to reject the predictions under this alternative specification.

[42] Bootstrapped values are estimated by drawing 1,000 samples of the data, each of which is constructed by drawing $m$ clusters of market-blocks with replacement, where $m$ is the number of original market-block clusters in the data.

unchanged, or if the entrant colludes. I observe that the mass of the distribution of the effects is only slightly to the right of what one would expect if the entrant colludes, and the 90% confidence interval can rule out alternative scenarios in which the entrant competes or even conduct remaining unchanged. This analysis, however, has used the point estimates of the demand parameters. Do we still have precision when taking into account the variance in these parameters? My analysis suggests we do. I solve for $\sigma_T$ in Equation 21:

$$\sigma_T = \frac{\delta(a - E[P_T])}{E[P_T - P_C] + \frac{\delta(a - E[P_C])}{\sigma_C}} \tag{23}$$

I then sample from the entire dataset 1,000 times. For each sample, I estimate $\rho$ and, using bootstrapped values of $\delta$, estimate $\sigma_C$. I then estimate $E[P_T]$, $E[P_C]$, and $E[P_T - P_C]$ for each sample. Finally, for each sample, I calculate $\sigma_T$ using Equation 23. The kernel density of the resulting $\sigma_C$ and $\sigma_T$ are displayed in Figure J.1. A Kolmogorov-Smirnov test cannot reject that these two distributions are the same (D=0.0183, p-val = 0.996). I therefore conclude that entry has left $\sigma$ unchanged. More specifically, I can test to which scenario the change in $\sigma$ most closely corresponds. Figure 10 presents these results, demonstrating that the mass of the change in $\sigma$ lies at zero, lining up closely with the predictions if entrants collude. The 90% confidence intervals rule out conduct remaining unchanged or the entrant competing.

### 7.4 Discussion

The limited observed change in price reveals that the entry generated had a negligible effect on competition. In an environment with a high degree of market power at baseline, this is consistent with entrants being able to easily join existing collusive agreements with incumbents upon arrival.[43] Further corroborating the interpretation of collusion is the fact that what little price decrease I observe is concentrated in markets with lower self-reports of collusion at baseline (see Figure 7).

The physical environment of these markets may contribute to the robustness of these agreements: traders can easily observe each others' transactions and in the absence of menu costs can quickly change prices if necessary to punish defectors. These features of the market layout may enable traders to collude even with new entrants who lack the history of repeated

[43] Note the importance of first understanding baseline levels of competition; without this, a null effect on price could be consistent with either a baseline market that is already perfectly competitive (and therefore entry would have no effect on price, which would already be at marginal costs) or a baseline market that is perfectly collusive in which the entrant joins in the collusive agreement.

interactions often required to maintain collusion in other settings.

Results here can of course only speak to the effect of entry in the context in which it was generated in this experiment. I document that traders generally exhibit low willingness to enter new markets, and those who did take up the offer tended to be larger and have more connections in the market. It may be that entry by different types of traders would have a greater effect on competition; however, the composition of entry seen in this experiment is likely the policy-relevant one, since larger and well-connected traders appear to be those that are responsive to nudges to encourage entry.

# 8 Conclusion

Policymakers have long speculated that agricultural traders in Africa exert market power, paying below-competitive prices to farmers and charging above-competitive prices to consumers. However, the absence of trader records and the difficulty in identifying clean shocks to traders' operating costs have challenged the ability of previous work to provide clear evidence on the nature of competition in this sector. In this paper, I present some of the first experimental evidence on the topic. I experimentally estimate pass-through and the curvature of demand, and use these parameters to calibrate a model of optimal pricing behavior. I find evidence of a high degree of intermediary market power. Welfare analysis suggests that consumers enjoy only 17.8% of the total variable surplus from these transactions, while intermediaries reap 67.6%. The remaining 14.6% is deadweight loss. In an additional experiment, I generate exogenous entry by offering traders subsidies to enter specific, randomly-selected markets in which they have never worked before. I find that each additional trader entering the market reduces prices by less than 1%. When interpreted through the lens of the model, this suggests that entrants collude with incumbents upon entry.

Taken together, these results suggests that policies commonly proposed to reduce the cost of agricultural trade – such as paving rural roads, implementing market price intelligence systems, and instituting uniform quality grading – would do little to achieve their stated aims of improving consumer and farmer welfare unless they also enhance the level of competition among traders.[44] Low pass-through rates indicate that traders retain the vast majority of reductions to their costs, rather than passing them on. Given the high degree of market power observed, policymakers may be interested in pursuing policies that explicitly target enhanced competition among intermediaries, which simulations indicate would yield large gains to consumers and improve market efficiency.

[44] Policies to reduce traders' costs may indirectly enhance competition (e.g., road construction could encourage entry into newly connected markets). However, I find here that entry does little to increase competition.

However, antitrust regulation of traders would likely be difficult to implement in an environment of low state capacity, and direct state intervention into the market to supplant the private sector would likely create more problems than it would solve, as seen during the largely disappointing experience with state-run markets following independence. Policies that encourage greater market entry may be more a feasible response. However, I find that entry yields little benefits to consumers, at least at levels seen in this experiment. While it is possible that massive entry – for example, doubling the number of traders in a market – could do more to increase competition (the effects of such a treatment are outside the scope of this paper), evidence presented here does suggest that such a policy is at best likely to be expensive, given that willingness to enter new markets appears low among most traders.

Identifying mechanisms that increase competition is therefore an open challenge, given that collusive agreements seem flexible in incorporating entrants. The physical layout of the market may contribute this flexibility. Selling directly next to each other, traders can easily observe each other's prices and readily respond to any deviations from agreement with a rapid price war. Further, consumers, who typically only shop in their local market, are captive to the traders there. More fundamental changes to the market environment may be needed to enhance competition.

New technologies, such as mobile marketplaces, hold some promise here. On these platforms, a larger pool of sellers interacts more anonymously, making coordination on price more difficult. Further, buyers can access a variety of sellers, rather than just those close to home. However, technological solutions must still address the real-world constraints of high transportation costs, limited trust, and other barriers that discourage exchange between new parties. The power of these technologies, as well as that of other potential mechanisms for expanding competition in these markets more broadly, is a ripe area for future research.

## References


**Aker, Jenny.** 2010. "Information from Market Near and Far: Mobile Phones and Agricultural Market in Niger." *American Economic Journal Applied Economics*, 2: 46–59.

**Allen, Treb.** 2014. "Information Frictions in Trade." *Econometrica*, 82(6): 2041–2083.

**Antras, Pol, and Arnaud Costinot.** 2011. "Intermediated Trade." *The Quarterly Journal of Economics*, 126: 1319–1374.

**Argent, Jonathan, and Tania Begazo.** 2015. "Competition in Kenyan Markets and Its Impact on Income and Poverty: A Case Study on Sugar and Maize." *World Bank Policy Research Working Ppaer*, 7179.

**Atkin, David, and Dave Donaldson.** 2015. "Who's Getting Globalized? The Size and Nature of International Trade Costs." *NBER Working Paper*, , (21439). Working Paper.

**Attanasio, Orazio, and Elena Pastorino.** 2015. "Nonlinear Pricing in Village Economies." *Working Paper.*

**Bardhan, Pranab, Dilip Mookherjee, and Masatoshi Tsumagari.** 2013. "Middlemen Margins and Globalization." *American Economic Journal: Microeconomics*, 5(4): 81–119.

**Berge, Lars Ivar Oppedal Berge, Kjetil Bjorvatn, Simon Galle, Edward Miguel, Daniel N Posner, Bergil Tungodden, and Kelly Zhang.** 2015. "How Strong are Ethnic Preferences?" *NBER Working Paper*, , (21715).

**Bulow, Jeremy I., and Paul Pfleiderer.** 1983. "A Note on the Effect of Cost Changes on Price." *Journal of Political Economy*, 91(1): 182–185.

**Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel.** 2016. "Selling Low and Buying High: An Arbitrage Puzzle in Kenyan Villages." *Working Paper.*

**Casaburi, Lorenzo, and Tristian Reed.** 2016. "Competition and Interlinkages in Agricultural Markets: An Experimental Approach." *Working Paper.*

**Casaburi, Lorenzo, Rachel Glennerster, and Tavneet Suri.** 2013. "Rural Roads and Intermediated Trade: Regression Discontinuity Evidence from Sierra Leone." *Working Paper.*

**Corts, Kenneth.** 1999. "Conduct Parameters and the Measurement of Market Power." *Journal of Econometrics*, 88: 227–250.

**Dillon, Brian.** 2016. "Selling Low to Pay for School: A Large-Scale Natural Experiment in Malawi." *Working Paper.*

**Dillon, Brian, and Chelsey Dambro.** 2016. "How Competitive are Food Crop Markets in Sub-Saharan Africa? A Review of the Evidence." *Working Paper.*

**Fafchamps, Marcel, and Bart Minten.** 2012. "Impact of SMS-Based Agricultural Information on Indian Farmers." *The World Bank Economic Review*, 26(3): 383–414.

**Fafchamps, Marcel, Eleni Gabre-Madhin, and Bart Minten.** 2005. "Increasing Returns and Market Efficiency in Agricultural Trade." *Journal of Development Economics*, 78: 406–442.

**Hildebrandt, Nicole, Yaw Nyarko, Giorgia Romagnoli, and Emilia Soldani.** 2015. "Price Information, Inter-Village Networks, and Bargaining Spillovers: Experimental Evidence from Ghana." *Working Paper.*

**Hjort, Jonas.** 2014. "Ethnic Divisions and Production in Firms." *Quarterly Journal of Economics*, 1899–1946.

**Jayne, T.S., Robert Myers, and James Nyoro.** 2008. "The Effects of NCPB Marketing Policies on Maize Market Prices." *Agricultural Economics*, 38: 313–325.

**Jensen, Robert.** 2007. "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector." *Quarterly Journal of Economics*, 3: 879–924.

**Keniston, Daniel E.** 2011. "Bargaining and Welfare: A Dynamic Structural Analysis." *Working Paper.*

**McKenzie, David, and Christopher Woodruff.** 2015. "Business Practices in Small Firms in Developing Countries." *Working Paper.*

**Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria.** 2015. "Asymmetric Information and Middleman Margins: An Experiment with Indian Potato Farmers." *HKUST IEMS Working Paper Series*, 29.

**Moser, Christine, Christopher Barret, and Bart Minten.** 2009. "Spatial Integration at Multiple Scales: Rice Markets in Madagascar." *Agricultural Economics*, 40: 281–294.

**Osborne, Theresa.** 2005. "Imperfect Competition in Agricultural Markets: Evidence from Ethiopia." *Journal of Development Economics*, 76: 405–428.

**Rashid, Shahidur, and Nicholas Minot.** 2010. "Are Staple Food Markets in Africa Efficient? Spatial Price Analysis and Beyond."

**Robinson, Amanda.** 2016. "Internal Borders: Ethnic-Based Market Segmentations in Malawi." *World Development*, 87: 371–384.

**Startz, Meredith.** 2017. "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade." *Working Paper.*

**Teravaninthorn, Supee, and Gael Raballand.** 2009. "Transport Prices and Costs in Africa." *World Bank.*

**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.

# Tables and Figures

Figure 1: **Maize prices in study markets**. Grey lines show the price for each market over the 12-week study period. The black line shows the average price across markets. The black bar on the vertical axis shows the average size of the cost reduction subsidy (2.22Ksh/kg).

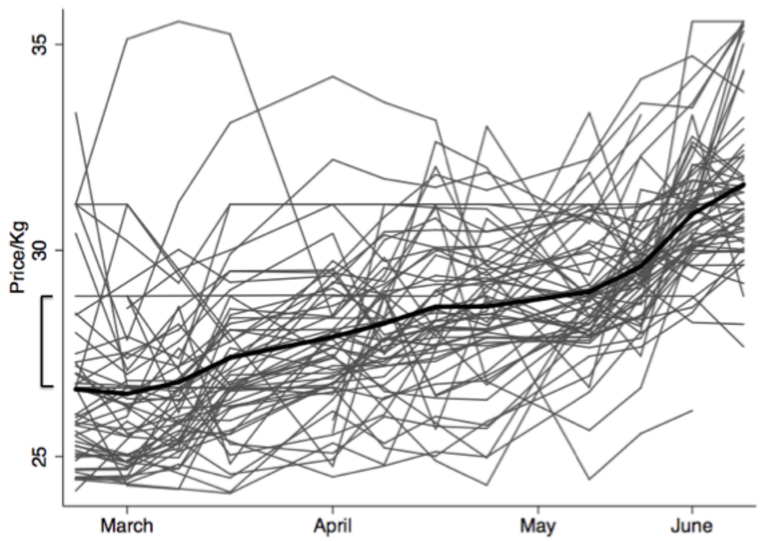


Figure 2: **Pass-through by market size**. Pass-through as estimated in markets of each size (bars represent the 95% confidence interval). The average for the full sample is 22% (dotted line). The bottom two estimates show pooled results, grouped into above/below median size.

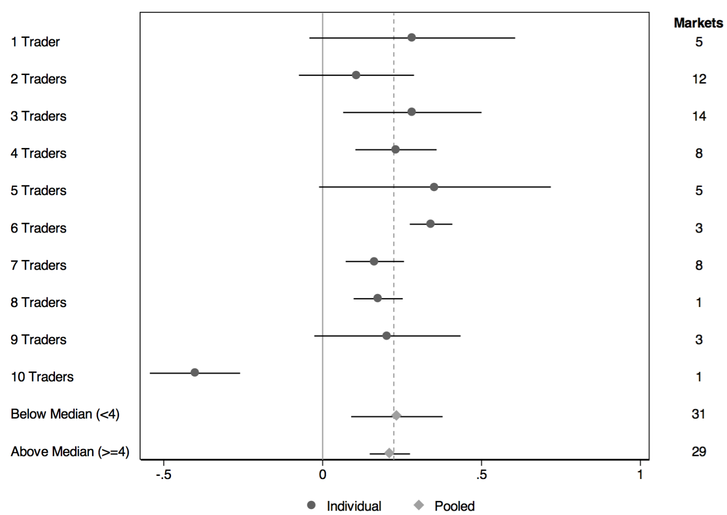

Figure 3: **Pass-through by various factors**. Pass-through as estimated in markets in each category (bars represent the 95% confidence interval). The average for the full sample is 22% (dotted line). Categories are: above/below median number of traders; on/off tarmac roads; above/below median number of days in which at least one trader reports discussing prices with other traders; above/below the median number of days in which at least one trader reports a price agreement.

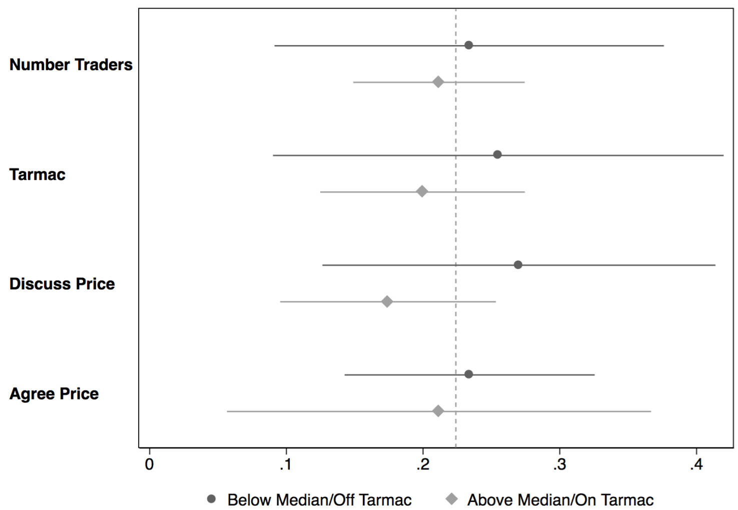


Figure 4: **Predicted pass-through under three models**. Given my demand curvature estimate, I predict that one would have observed 100% pass-through in a perfectly competitive market, 55% pass-through in a Cournot competitive market, and 26% pass-through in a collusive market environment. The distribution of empirical pass-through, calculated for 1,000 bootstrapped samples, is shown in grey. The point estimate and 90% confidence interval is shown in red.

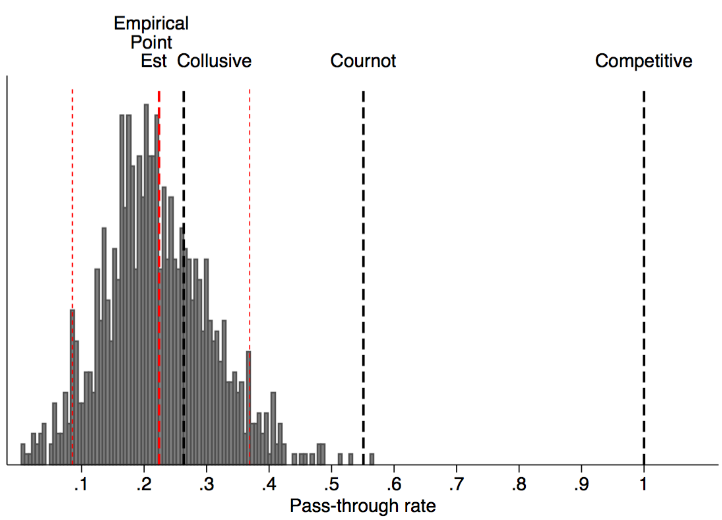

Figure 5: **Competitiveness parameter (sigma) estimates**. 1,000 bootstrapped estimates of $\sigma$ are identified by calibrating Equation 14 with 1,000 bootstrapped estimates of $\rho$ and $\delta$. Recall that $\sigma = 1$ if competitive, $N$ if Cournot competitive, and $\infty$ if perfectly competitive. Because $\sigma = N$ varies across markets, for the Cournot benchmark I show the average $\sigma$ one should expect to see if all markets are behaving in a Cournot competitive manner, given the distribution of the market sizes in my sample. The competitive benchmark of $\sigma = \infty$ is not shown here for obvious reasons.

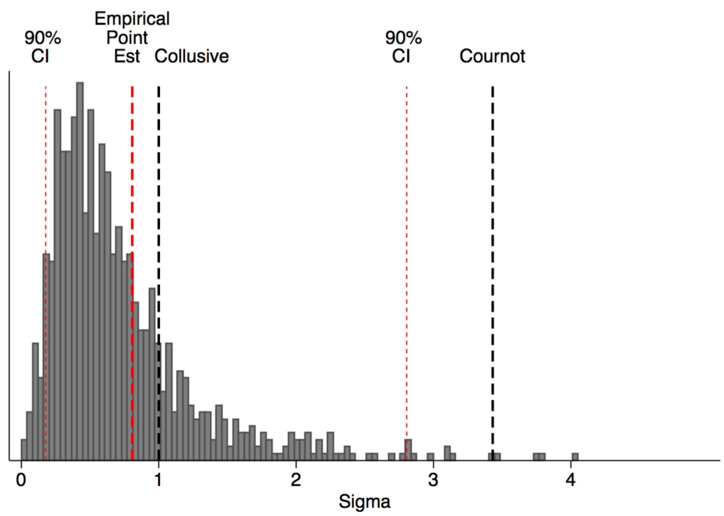


Figure 6: **Welfare counterfactuals**. Counterfactual division of welfare is shown for the average market size of four traders. The current division of surplus is shown at the far left vertical dotted line, suggesting that intermediary surplus (IS) is 67.6% of total surplus, while consumer surplus (CS) is only 17.6% and deadweight loss (DWL) is 14.6%. Movements to the right represent increases in competition. Dotted vertical lines at "Cournot" and "competitive" identify how this division would be altered if the market operated under these models of enhanced competition.

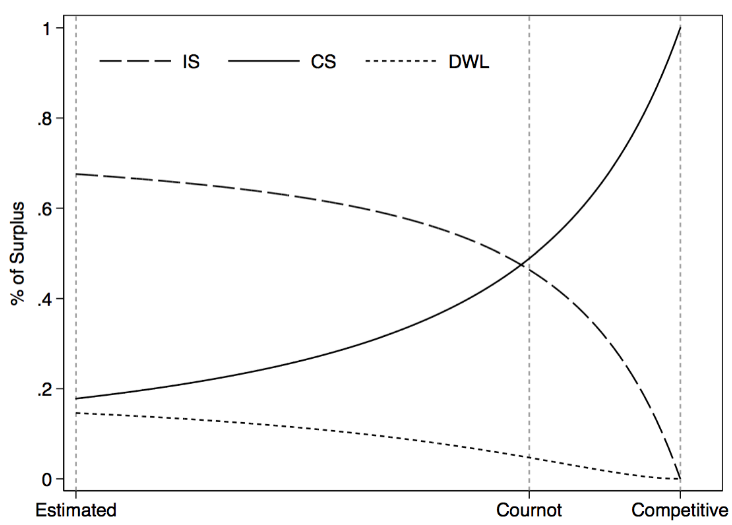

Figure 7: **Heterogeneity in willingness-to-enter**. Take-up of the entry offer regressed on various measures of heterogeneity (alternately without and with market fixed effects; the latter compares only traders offered to attend the same market). The coefficient and 95% confidence interval is plotted.

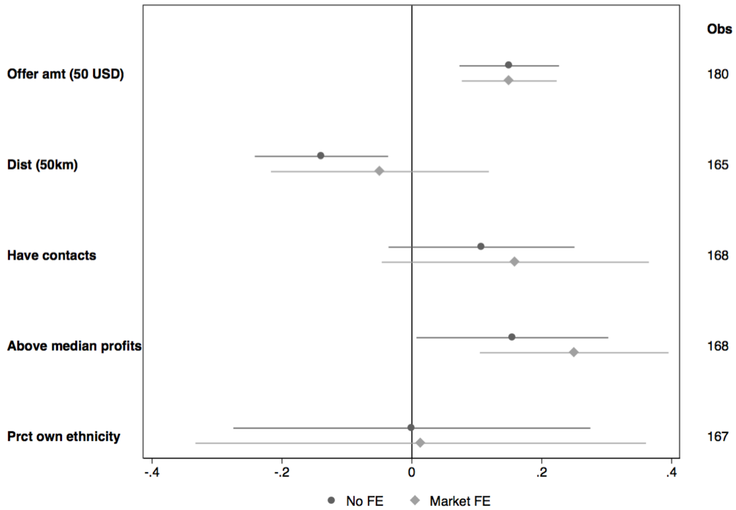


Figure 8: **Heterogeneity in take-up and IV impact of entry by market characteristics**. Categories are: above/below median number of traders; on/off tarmac roads; above/below median number of days in which at least one trader reports discussing prices with other traders; above/below the median number of days in which at least one trader reports a price agreement. The unit of observation in Panel A is the market-day and the sample is restricted to entry treatment market days. Panel A presents the results from a t-test of a dummy for whether any entry occurred on that market-day by the relevant dummy. The mean and 95% confidence intervals for each subgroup are shown. Panel B uses the full trader sample and presents the point estimate and standard errors on an IV specification identical to that presented in Equation 18, but with the sample restricted to the subgroup in question.

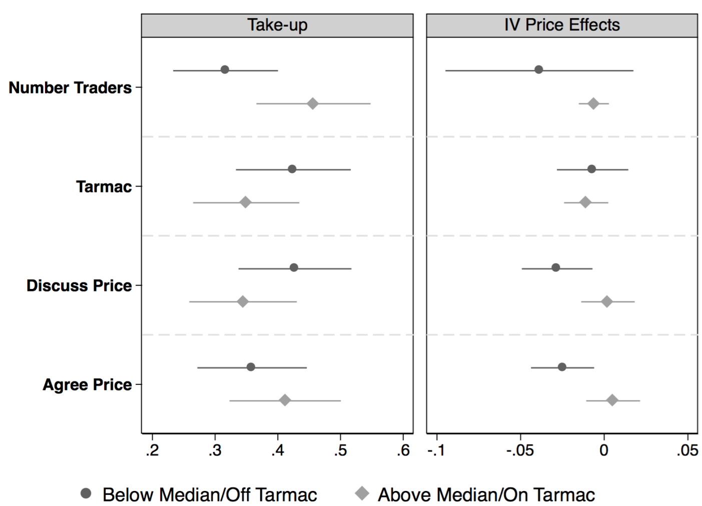

Figure 9: **Predicted price change for three models of entrant behavior**. Given my demand parameter estimates, the distribution of market sizes, and the first-stage effect of increasing N by 0.582, I predict that one would have observed a 7% reduced form reduction in prices if the entrant were to compete with incumbents, a 4% reduction if market conduct were unchanged upon entry, and 0% reduction if the entrant were to collude with incumbents. The distribution of empirical price effects, calculated for 1,000 bootstrapped samples, is shown in grey. The point estimate and 90% confidence interval is shown in red.

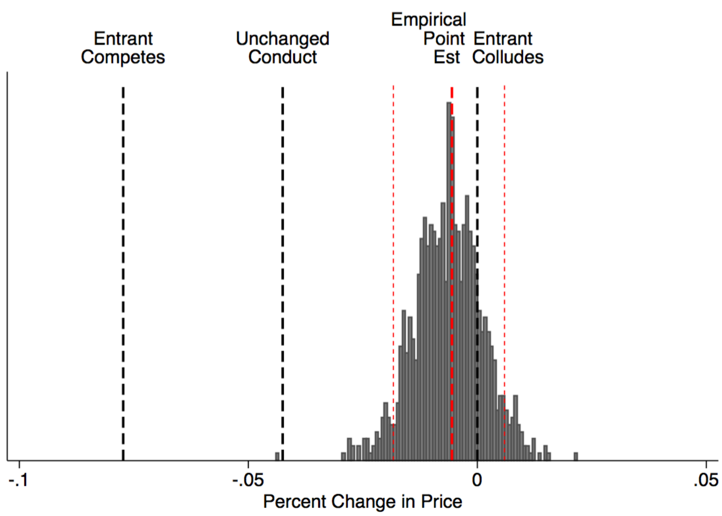


Figure 10: **Distribution of change in sigma**. The distribution of the change in the competitiveness parameter ($\sigma_T - \sigma_C$) from 1,000 bootstrapped samples is plotted. The 90% confidence interval is shown in red dotted lines. The vertical black dotted lines present the expected change in the competitiveness parameter under the three models considered.

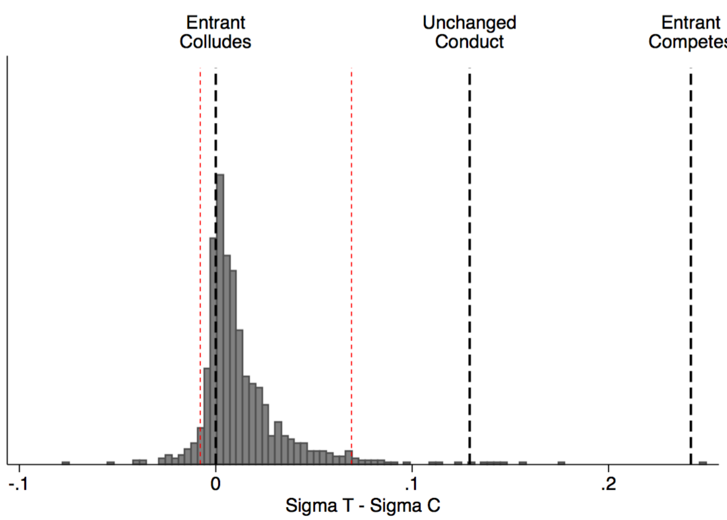

Table 1: **Pass-through.** The first column shows the overall pass-through rate of 22%. The second column shows pass-through rates separately by "low" and "high" offers.

| | (1) Price | (2) Price |
|---|---|---|
| Cost Reduction | 0.224 | |
| | (0.0434) | |
| Cost Reduction - Low | | 0.219 |
| | | (0.0538) |
| Cost Reduction - High | | 0.228 |
| | | (0.0618) |
| Mean Dep Var | 28.92 | 28.92 |
| N | 1860 | 1860 |
| Market FE | Yes | Yes |
| Week FE | Yes | Yes |

Table 2: **Demand Estimation.** The point estimates and the 95% confidence intervals for the three estimated parameters of the Bulow-Pfleiderer demand function are displayed

| | **Parameter Estimate** | **Lower Bound** | **Upper Bound** |
|---|---|---|---|
| a | 42.50 | 42.22 | 57.60 |
| b | 0.0006 | 0.0000 | 0.0671 |
| $\delta$ | 2.80 | 0.01 | 5.89 |

Table 3: **Welfare Estimates.** The first row shows consumer surplus, intermediate surplus, and deadweight loss at the closest theory-consistent $\sigma$ of 1 (and the corresponding pass-through rate of 26%). The second row presents the upper 95% confidence interval estimates of consumer welfare, which is maximized when using the upper end of the confidence interval on $\rho$ and $\sigma$.

| | **Consumer Surplus** | **Intermediary Surplus** | **DWL** |
|---|---|---|---|
| Point Estimate | 0.178 | 0.676 | 0.146 |
| Upper 95% CI on CS | 0.255 | 0.661 | 0.084 |

Table 4: **Welfare Counterfactuals.** The first row shows point estimates for consumer surplus, intermediate surplus, and deadweight loss at the closest theory-consistent $\sigma$ of 1 (and the corresponding pass-through rate of 26%). The second row presents the counterfactual welfare distribution if markets were Cournot competitive for the average market of four traders. The third row presents the counterfactual welfare if markets were perfectly competitive.

| | **Consumer Surplus** | **Intermediary Surplus** | **DWL** |
|---|---|---|---|
| Current Environment | 0.178 | 0.676 | 0.146 |
| Cournot Competitive | 0.489 | 0.464 | 0.047 |
| Perfectly Competitive | 1.000 | 0.000 | 0.000 |

Table 5: **Take-up of Entry Offers.** Offers ranged from 5,000-15,000 Kenyan shillings ($49-148 USD). "Take-up" = 1 if the trader *ever* took up an offer during any of the four weeks for which the offer was available. "Report profitable"=1 when the profits expected under take-up + the offer amount > profits expected under no take-up (as repaired by traders during the offer phone call).

| | Offer Amount | | Take-up | Report | Obs |
|---|---|---|---|---|---|
| | *Ksh* | *USD* | Rate | Profitable | |
| **Low Offer** | 5,000 | 49 | 0.12 | 0.77 | 60 |
| **Medium Offer** | 10,000 | 99 | 0.28 | 0.80 | 60 |
| **High Offer** | 15,000 | 148 | 0.42 | 0.89 | 60 |

Table 6: **Effect of Entry.** The variable "Entry Offer Market" is a dummy for treatment status in the entry experiment. "Num Traders" is the number of traders present in the market on that day. Column 1 presents the first stage effect of treatment on the number of traders. Column 2 presents the reduced form effect of treatment on log price. Column 3 presents the effect of the number of traders on the log price, instrumenting for the number of traders with treatment.

| | (1) Num Traders | (2) Ln Price | (3) Ln Price |
|---|---|---|---|
| Entry Offer Market | 0.582 | -0.00555 | |
| | (0.118) | (0.00357) | |
| Num Traders | | | -0.00955 |
| | | | (0.00582) |
| Type | FS | RF | IV |
| F-Stat FS | | | 24.42 |
| Mean Dep Var | 4.427 | 3.364 | 3.364 |
| N | 1776 | 1776 | 1776 |
| Market FE | Yes | Yes | Yes |
| Week FE | Yes | Yes | Yes |

Table 7: **Simulated effect of entry on prices for various models of entrant behavior.** Estimated effect on $\theta$, $\sigma$, and prices under various forms of entrant behavior. Effects are identified separately at each level of baseline number of traders in the market. Initial baseline $\sigma_C$ is assumed to be 1 and demand parameters are taken at their point estimate. The **top panel** presents simulated IV effects of impact of entry by one trader. The **middle panel** presents simulated reduced form effects of the impact of entry, given the first-stage increase in the number of traders observed at each market-size. Because these first-stage effects may contain substantial noise, given the small cells of some market-size buckets, I present in the **bottom panel** simulated reduced form effects of the impact of entry using the average first-stage for all market-sizes.

| **Baseline Num Traders** | **Num Mkts** | **$\Delta$ N** | **Conduct Unchanged** | | | **Entrant Competes** | | | **Entrant Colludes** | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | *Theta* | *Sigma* | *% Price $\Delta$* | *Theta* | *Sigma* | *% Price $\Delta$* | *Theta* | *Sigma* | *% Price $\Delta$* |
| ***IV*** | | | | | | | | | | | |
| 1 Trader | 5 | 1.00 | 1 | 2.00 | -26.98 | 1.00 | 2.00 | -26.98 | 1.00 | 1 | 0 |
| 2 Traders | 12 | 1.00 | 2 | 1.50 | -15.06 | 1.67 | 1.80 | -22.53 | 0.50 | 1 | 0 |
| 3 Traders | 14 | 1.00 | 3 | 1.33 | -10.45 | 2.50 | 1.60 | -17.66 | 0.33 | 1 | 0 |
| 4 Traders | 8 | 1.00 | 4 | 1.25 | -8.00 | 3.40 | 1.47 | -14.27 | 0.25 | 1 | 0 |
| 5 Traders | 5 | 1.00 | 5 | 1.20 | -6.48 | 4.33 | 1.38 | -11.91 | 0.20 | 1 | 0 |
| 6 Traders | 3 | 1.00 | 6 | 1.17 | -5.44 | 5.29 | 1.32 | -10.19 | 0.17 | 1 | 0 |
| 7 Traders | 8 | 1.00 | 7 | 1.14 | -4.69 | 6.25 | 1.28 | -8.89 | 0.14 | 1 | 0 |
| 8 Traders | 1 | 1.00 | 8 | 1.12 | -4.13 | 7.22 | 1.25 | -7.88 | 0.12 | 1 | 0 |
| 9 Traders | 3 | 1.00 | 9 | 1.11 | -3.68 | 8.20 | 1.22 | -7.07 | 0.11 | 1 | 0 |
| 10 Traders | 1 | 1.00 | 10 | 1.10 | -3.32 | 9.18 | 1.20 | -6.42 | 0.10 | 1 | 0 |
| ***By Market RF*** | | | | | | | | | | | |
| 1 Trader | 5 | -0.02 | 1 | 0.98 | 0.73 | 1.00 | 0.98 | 0.73 | 1.00 | 1 | 0 |
| 2 Traders | 12 | 0.28 | 2 | 1.14 | -4.65 | 1.88 | 1.22 | -7.00 | 0.50 | 1 | 0 |
| 3 Traders | 14 | 0.29 | 3 | 1.10 | -3.27 | 2.82 | 1.17 | -5.49 | 0.33 | 1 | 0 |
| 4 Traders | 8 | 0.80 | 4 | 1.20 | -6.45 | 3.50 | 1.37 | -11.48 | 0.25 | 1 | 0 |
| 5 Traders | 5 | 1.08 | 5 | 1.22 | -6.99 | 4.29 | 1.42 | -12.85 | 0.20 | 1 | 0 |
| 6 Traders | 3 | 0.91 | 6 | 1.15 | -4.97 | 5.34 | 1.29 | -9.29 | 0.17 | 1 | 0 |
| 7 Traders | 8 | 0.20 | 7 | 1.03 | -0.99 | 6.83 | 1.06 | -1.85 | 0.14 | 1 | 0 |
| 8 Traders | 1 | 1.76 | 8 | 1.22 | -7.09 | 6.74 | 1.45 | -13.68 | 0.12 | 1 | 0 |
| 9 Traders | 3 | 1.08 | 9 | 1.12 | -3.97 | 8.14 | 1.24 | -7.63 | 0.11 | 1 | 0 |
| 10 Traders | 1 | 8.43 | 10 | 1.84 | -23.50 | 5.88 | 3.13 | -46.53 | 0.10 | 1 | 0 |
| ***Average RF*** | | | | | | | | | | | |
| 1 Trader | 5 | 0.58 | 1 | 1.58 | -17.20 | 1.00 | 1.58 | -17.20 | 1.00 | 1 | 0 |
| 2 Traders | 12 | 0.58 | 2 | 1.29 | -9.21 | 1.77 | 1.45 | -13.85 | 0.50 | 1 | 0 |
| 3 Traders | 14 | 0.58 | 3 | 1.19 | -6.29 | 2.68 | 1.34 | -10.61 | 0.33 | 1 | 0 |
| 4 Traders | 8 | 0.58 | 4 | 1.15 | -4.78 | 3.62 | 1.27 | -8.48 | 0.25 | 1 | 0 |
| 5 Traders | 5 | 0.58 | 5 | 1.12 | -3.85 | 4.58 | 1.22 | -7.03 | 0.20 | 1 | 0 |
| 6 Traders | 3 | 0.58 | 6 | 1.10 | -3.22 | 5.56 | 1.18 | -5.99 | 0.17 | 1 | 0 |
| 7 Traders | 8 | 0.58 | 7 | 1.08 | -2.77 | 6.54 | 1.16 | -5.22 | 0.14 | 1 | 0 |
| 8 Traders | 1 | 0.58 | 8 | 1.07 | -2.43 | 7.53 | 1.14 | -4.62 | 0.12 | 1 | 0 |
| 9 Traders | 3 | 0.58 | 9 | 1.06 | -2.17 | 8.51 | 1.13 | -4.14 | 0.11 | 1 | 0 |
| 10 Traders | 1 | 0.58 | 10 | 1.06 | -1.95 | 9.51 | 1.11 | -3.75 | 0.10 | 1 | 0 |

## For Online Publication

# A Appendix: Maize Value Chains and Trader 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 my 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).

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

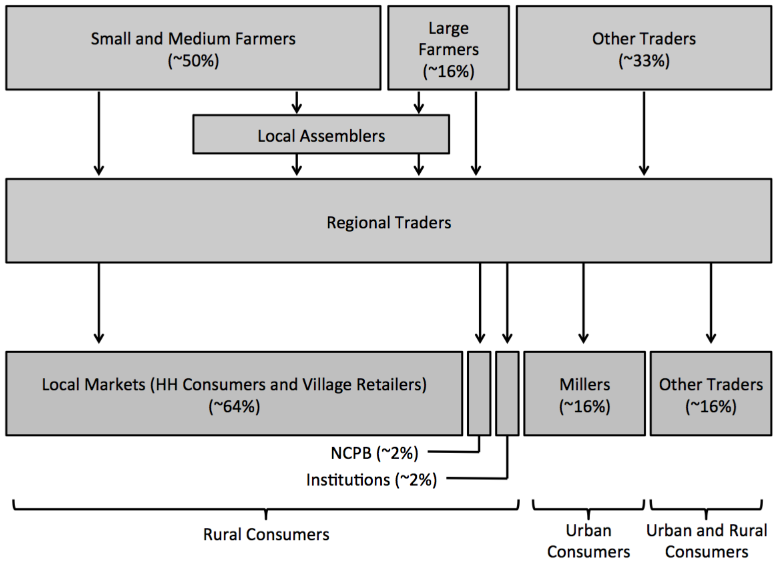


Table A.1: **Trader summary statistics.**

| | 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 | 0.95 | 0.22 | 2,964 |
| New trader | 0.02 | 0.13 | 2,964 |
| Worked with all before | 0.77 | 0.42 | 3,038 |
| Know other traders well | 0.67 | 0.47 | 2,571 |
| Know other traders well or somewhat well | 0.94 | 0.24 | 2,571 |
| ***Collusion Reports*** | | | |
| Self-report discuss price | 0.38 | 0.49 | 2,571 |
| Someone in market report discuss price | 0.80 | 0.40 | 2,806 |
| Percent traders with whom discuss price | 0.77 | 0.28 | 977 |
| Self-report agree price | 0.30 | 0.46 | 2,571 |
| Someone in market report agree price | 0.72 | 0.45 | 2,806 |
| Percent traders with whom agree price | 0.77 | 0.28 | 778 |

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

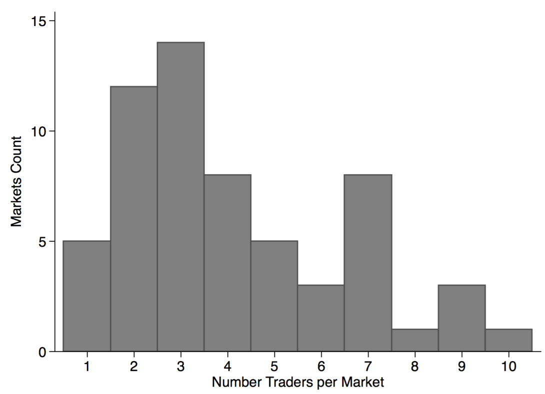

# B Appendix: Static Model

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 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 I observe. Related work in the region suggests that credit constraints limit households' ability to arbitrage these price fluctuations (Burke, Bergquist and Miguel, 2016).

The randomized order of treatment periods allows me 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, I 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 B.1 presents the results for the full sample. I see that having been a pass-through treated market in the previous 4-week block does not affect the quantities sold 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. I continue to see no evidence of a stockpiling effect here (in fact, the point estimate becomes positive, though standard errors also increase substantially with this reduced sample). Given limited evidence of consumer stockpiling, I model demand as static and therefore decisions regarding prices and quantities as separable across market-days.

Table B.1: **Effect of Previous Treatment Status on Quantity Sold Today.** Log quantity sold as a function of previous treatment status, controlling for current treatment status. "PT Previous" is a dummy for whether the market was in a pass-through treatment market in the previous period. Column 1 presents results for the full sample. Column 2 presents results for the first week of the block, when one would expect to see most concentrated dynamic effects, if existent.

| | Ln Kgs | Ln Kgs |
|---|---|---|
| PT Previous | -0.0131 | 0.199 |
| | (0.157) | (0.213) |
| Mean DV | 7.369 | 7.273 |
| N | 2191 | 541 |
| Sample | Full Block | Week 1 Only |
| Market FE | Yes | Yes |
| Week FE | Yes | Yes |

## C Appendix: Product Differentiation

Staple food commodities are often pointed to as the textbook example of a homogenous goods. However, I 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. I 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, I 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].[45]

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.[46] I therefore collect only one measurement of quality for each trader in each market-day. Across traders in the same market day I 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.1, 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.[47] Moreover, while I do see small price differences for purchases on credit, this relationship disappears when controlling for other features of the transaction.[48] Therefore, the weight

[45] 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.

[46] 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.

[47] 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 (results available upon request).

[48] Unexpectedly, the relationship between credit and price seen in Column 2 is negative, but this may be

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

Table C.1: **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)<br>Ln Price | (2)<br>Ln Price | (3)<br>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 |

Despite the weight of empirical evidence against meaningful product differentiation, it is worth noting that — theoretically — the prediction for pass-through under perfect collusion as outlined in Section 2 is observationally equivalent to that under an alternative market structure in which traders sell perfectly differentiated products (i.e., when consumers' elasticity of substitution across products is zero). In this alternate structure, one could model traders as monopolists working in their own "markets" with an $N = 1$ and a $\theta = 1$ (and therefore $\sigma = 1$) as they are the only trader selling that particular type of good. If one assumes that $E$ is the same across each traders' segment of consumers (a nontrivial assumption, but perhaps a reasonable first approximation as the elasticity of the slope of inverse demand is invariant to the size of the market), the pass-through rate would be the same as under perfect collusion, since $\sigma$ and $E$ would be identical under the two scenarios. Therefore, a pass-through rate consistent with market power from perfect collusion is also consistent with market power from perfect product differentiation; I cannot theoretically distinguish between the two. Given the strong empirical evidence that maize is a homogeneous good, I interpret market power as arising from collusion. Moreover, whether exerted by collusion or differentiation, the low pass-through rate is evidence of a high degree of market power.

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).

# D Appendix: Price Discrimination

Empirically, I 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).[49] To explore whether there is evidence of such nonlinear pricing schemes in my setting, I utilize transaction-level data (totaling 39,667 transactions) and explore the covariance of price and quantity of maize sold by the same trader to his customers in a given market-day. Figure D.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, I 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, which the bulk of overall variation of price lies 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 D.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.

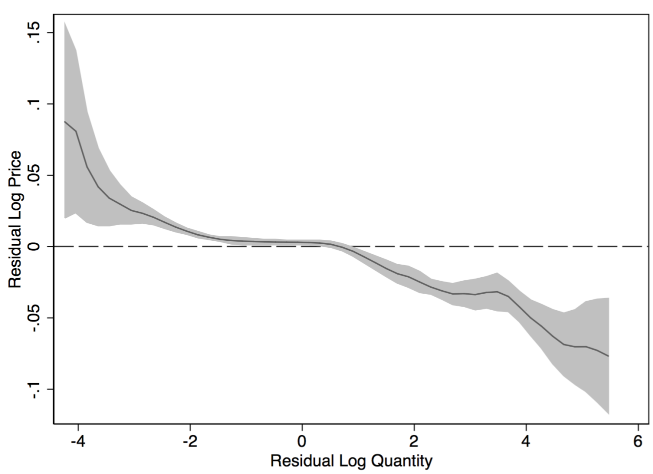


[49]Attanasio and Pastorino (2015) find that sellers of food staples in Mexico are able to exert market power to discriminate across customers with different levels of willingness (and ability) to pay. Sellers in their setting offer nonlinear pricing schemes using bulk discounts.

# E Appendix: Constant Marginal Costs

A key assumption of the model underpinning this exercise is that of constant marginal costs. In this appendix, I (i) discuss why this assumption is important; (ii) present empirical evidence supporting this assumption; and (iii) show that even if this assumption were relaxed, an implausibly steep increasing marginal costs curve would be necessary to reconcile the observed pass-through rate with perfect competition or Cournot competition, the two alternative models rejected by the analysis in the main text. The primary conclusions of this paper are therefore unlikely to be strongly driven by this assumption.

## E.1 The Constant Marginal Cost Assumption

It is worth highlighting why this assumption is an important one. To do so, first note that in a more general model without the assumption of constant marginal cost, one needs to differentiate between the pass-through of a marginal costs shock $\rho = \frac{\partial P}{\partial c}$ and the pass-through of the cost reduction subsidy $\tilde{\rho} = \frac{\partial P}{\partial CR}$. Under the assumption of constant marginal costs, the two are equivalent, as the subsidy simply represents a vertical shift down in a flat marginal cost curve. However, when marginal costs vary with quantity, the subsidy not only induces a downward shift in the cost curve, but also a shift *along* this curve, as quantity expands. As a result, depending on whether marginal costs are increasing or decreasing in quantity, marginal costs will respectively adjust by less or more than the amount of the subsidy. The concern, then, is that under a model of increasing costs, a low subsidy pass-through rate may simply reflect a small shift in marginal costs, rather than low cost pass-through.

Theory helps clarify and quantify this concern. In the main text, I developed a model for $\rho$, predicted what $\rho$ one should expect to observe empirically under three different forms of competition, and compared the observed $\rho$ to these three benchmarks in order to identify the model that best describes the empirical context. However, what I actually observe empirically – the 22% pass-through estimated in Equation 10 – is the impact of a one-unit increase in the cost reduction subsidy on price, or $\tilde{\rho}$ (which may or may not be equal to $\rho$ under a more relaxed assumption regarding the marginal cost curve). To align the theory to this relaxed assumption, I must develop a model of $\tilde{\rho}$.

I start from a slightly adjusted equation for profits, which can no longer can be written as $\pi_{dt} = (P_{dt} - c_{dt})q_{dt}$ (as in Equation 1) when marginal costs may not be constant. Now, firm profits must be written as total revenue - total costs (- any cost reduction subsidy):

$$\pi_{dt} = P_{dt}q_{dt} - TC_{dt} - CR_{dt}q_{dt} \tag{24}$$

where $CR_{dt}$ is the *negative* value of the marginal subsidy in pass-through treatment markets and zero elsewhere (as in Section 4). Taking the derivative of Equation 24 with respect to quantity $q_{dt}$ yields the trader's first order condition:

$$P_{dt} = \frac{\partial TC_{dt}}{\partial q_{dt}} + CR_{dt} - \theta \frac{\partial P_{dt}}{\partial Q_{dt}} \frac{Q_{dt}}{N_{dt}} = \frac{\partial TC_{dt}}{\partial q_{dt}} + CR_{dt} - \frac{1}{\sigma} \frac{\partial P_{dt}}{\partial Q_{dt}} Q_{dt} \tag{25}$$

In order to get an estimate of $\tilde{\rho} = \frac{\partial P}{\partial CR}$ the derivative of Equation 25 with respect to $CR_{dt}$ yields:

(26)
$$\tilde{\rho_{dt}} \equiv \frac{\partial P_{dt}}{\partial CR_{dt}} = \frac{\partial^2 TC_{dt}}{\partial q_{dt}^2}\frac{1}{N}\frac{\partial Q_{dt}}{P_{dt}}\frac{\partial P_{dt}}{CR_{dt}} + 1 - \frac{1}{\sigma}\{\frac{\partial^2 P_{dt}}{\partial Q_{dt}^2}\frac{\partial Q_{dt}}{\partial P_{dt}}\frac{\partial P_{dt}}{\partial CS_{dt}}Q_{dt} + \frac{\partial P_{dt}}{\partial Q_{dt}}\frac{\partial Q_{dt}}{\partial P_{dt}}\frac{\partial P_{dt}}{\partial CS_{dt}}\}$$

Which simplifies to

$$\tilde{\rho_{dt}} = \left\{1 + \frac{1+E_{dt}}{\sigma_{dt}} - \frac{\partial^2 TC_{dt}}{\partial q_{dt}^2}\frac{1}{N}\frac{\partial Q_{dt}}{P_{dt}}\right\}^{-1} \tag{27}$$

Note that the first portion of Equation 27 is identical to the original Equation 6. Only the last term differs. Therefore the key to identifying how $\tilde{\rho}$ differs from the originally estimated pass-through rate is the sign of $\frac{\partial^2 TC}{\partial q^2}$; that is, are marginal costs constant, increasing, or decreasing? In Section 2, I assume that marginal costs are are equal to a constant $c$ (that is, that total costs take the form $TC = cq + F$). In this case, $\frac{\partial^2 TC}{\partial q^2} = 0$ and Equation 27 reduces to the original Equation 6. However, if $\frac{\partial^2 TC}{\partial q^2} > 0$ (that is, if marginal costs are increasing), then the predicted pass-through rate for any given model of competition will be lower than under the assumption of constant marginal costs. Conversely, if $\frac{\partial^2 TC}{\partial q^2} < 0$ (if marginal costs are decreasing), the predicted pass-through rate for any given model of competition will be higher.

### E.2 Empirical Evidence for Constant Marginal Costs

I turn now to the empirical evidence regarding the sign of $\frac{\partial^2 TC}{\partial q^2}$. This evidence, limited though it may be, suggests that the assumption of $\frac{\partial^2 TC}{\partial q^2} = 0$ 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, Gabre-Madhin and Minten, 2005). Despite the difficulties I have noted surrounding such direct cost estimates,[50] I also try to ask about some of these costs in my setting, and find they look appear to be fairly constant with respect to quantity (see Table

[50] Estimating costs directly are made difficult by the fact that only 58% of traders keep any written records, and those that do typically keep only rudimentary records on sales, omitting most cost categories. Own-labor, a major input, is difficult to price. Finally, self-reported costs are subject to potential self-reporting bias. Traders are widely perceived as exploitative middlemen who rip off both the farmers from whom they

E.1). Therefore, the available empirical evidence suggests that constant marginal costs is a reasonable approximation of the empirical setting in which this experiment takes place.

Table E.1: **Costs as a Function of Quantity Sold.** The dependent variable in Column 1 is the per-kg cost of the purchasing inventory (likely the largest variable cost that traders face). The dependent variable in Column 2 is a rough proxy for "total cost" per-kg, estimated by asking at cost-per-kg traders would need to charge per-kg to break-even. Simply regressing these costs on quantity sold may produce biased estimates of the cost curve, as many factors may be correlated with operating a larger-volume business. The pass-through experiment, however, nicely offers an instrument for increased volume of sales. I therefore use treatment status as an instrument for kilograms sold. Costs do not appear to vary with quantity sold. Point estimates are fairly precise zeros, insignificant both in terms of statistical precision and magnitude.

| | (1)<br>Purchase Price | (2)<br>Total Costs |
|---|---|---|
| Kgs | -0.000584 | -0.000266 |
| | (0.000408) | (0.000350) |
| F-stat | 20.34 | 21.83 |
| N | 1889 | 1857 |
| Market FE | Yes | Yes |
| Week FE | Yes | Yes |

### E.3 Bounding the Impacts of the Constant Marginal Costs Assumption

Still, because this is an important assumption, and because I am not able to precisely pin down the slope of the cost curve, I here explore the implications of relaxing this assumption. Recall that the sign of $\frac{\partial^2 TC}{\partial q^2}$ is the key determinant of whether $\tilde{\rho}$ will be greater or less than originally estimated. While in theory the cost curve could take any form with respect to quantity, in order to discipline this exercise, I assume that marginal costs are linear in quantity, such that $TC = \frac{\alpha q^2}{2} + cq + F$ and $\frac{\partial^2 TC}{\partial q^2 = \alpha}$. Equation 27 now simplifies to:

$$\tilde{\rho_{dt}} = \left\{ 1 + \frac{1 + E_{dt}}{\sigma_{dt}} - \frac{\alpha}{N} \frac{\partial Q_{dt}}{P_{dt}} \right\}^{-1} \tag{28}$$

I now explore relaxing the assumption that $\alpha = 0$ and ask what $\alpha$ would be required to reconcile the observed $\tilde{\rho}$ of 22% with models other than the collusive model. Put another way, how wrong does the assumption of marginal costs have to be to alter the conclusion of this paper that collusion is at play?

purchase and the consumers to whom they sell. While the price at which they sell is directly observable to the researcher, the price they pay to farmers is not, as these purchases happen primarily at hard-to-observe geographically disburse farm-gate locations (note that this is a major benefit of the methodological approach used in this paper, which relies only on access to the observable sale price). Traders therefore may face social pressure to inflate the price they report paying to farmers in order to minimize their perceived mark-up.

In order to identify the required $\alpha$, I calibrate Equation 28 with the observed $\tilde{\rho}$ of 0.22 and $\frac{\partial Q}{\partial P} = 0.95$ (arrived at by evaluating $\frac{\partial Q}{\partial P}$ at median price of 29 Ksh/kg using the demand parameters as estimated in Section 5). According to the definition of $\sigma$, the competitiveness parameter, I calibrate the Cournot model with a $\sigma$ of $N = 4$ (recall that $\sigma = N$ in Cournot markets, with four being the median number of traders per market) and the competitive model with a $\sigma$ of $\infty$. I estimate that an $\alpha$ of 12.07 would be required to reconcile the observed pass-through rate with a Cournot model with increasing marginal costs and an $\alpha$ of 15.02 to reconcile with a competitive model. Both represent implausibly large slopes, with rapidly increasing marginal costs per kg. Further, the results in Table E.1 suggest that a slope of this magnitude is unlikely; both 12.07 and 15.02 lie well outside the 95% confidence interval. It is therefore difficult to reconcile the observed pass-through rate with alternate models by relaxing the assumption of constant marginal costs.

# F Appendix: Demand Curvature

Here I present a visual example of identical pass-through rates generated by two different underlying models: one in which markets are collusive and demand is only somewhat curved (left panel) and one in which markets are Cournot competitive and demand is strongly curved (right panel). In the left panel, I present a collusive environment, in which traders act as a single profit-maximizing firm. Quantities are set where the marginal revenue (MR) curve meets the marginal cost (MC) curve. Prices are then set where this optimal quantity intersects the demand (D) curve. A shift in the marginal cost curve downwards results in a pass-through rate of $\frac{\Delta P}{\Delta MC}$. With a concave demand function, as shown here, pass-through rates will be low.

In the right panel, I show a Cournot competitive environment, in which traders compete on quantities. The figure shows an individual firm's pricing and quantity decision. The firm takes the amount produced by other firms in the market ($q'$) as given (producing a residual demand (RD) curve) and from this determines its marginal revenue curve. It then sets its own quantities where its marginal cost curve meets its marginal revenue curve. Prices are then set by where total quantities hit the demand curve. This combination of competitive environment and demand curvature yields an identical pass-through rate to that seen in the left panel, despite coming from a different underlying model.

Figure F.1: **Pass-through given conduct and demand curvature**. **Left panel**: Collusion and slight curved demand. **Right panel**: Cournot competition and very curved demand.

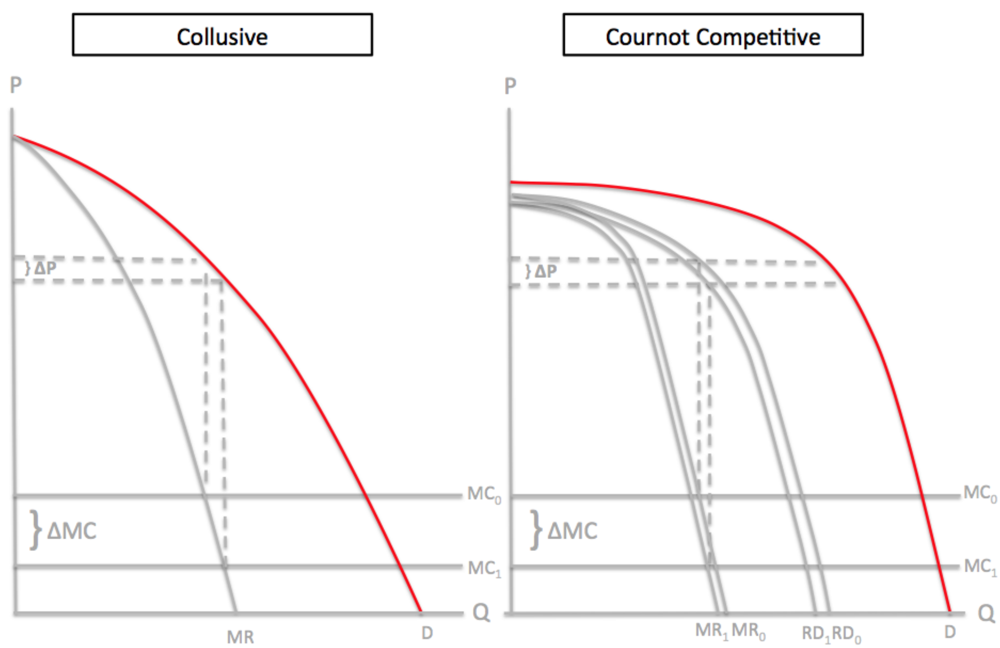

# G Appendix: Sample Selection and Experimental Schedule

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. I excluded markets that were reported to typically not have any maize traders 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.[51]

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.[52] I administered a pre-experiment survey of to this group of 60 selected markets in which I verified information provided by the Director of Trade and recorded the number of traders typically in the market.[53]. In a large number of these markets, it was found that the information provided by the Director of Trade was inaccurate.[54] Markets that were deemed ineligible upon visit were then replaced with market from their same stratum.[55] 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 G.1 presents the experimental schedule. The 60 markets in my 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 my 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.

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

[52]The 40 markets within 50km of Bungoma were selected randomly. This randomization was stratified to include 25 markets from which I 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 (in the case of a tie, one is randomly dropped).

[53]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.

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

[55]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 G.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>

## H Appendix: How do Entrants Compare to Incumbents?

This appendix provides greater detail on how incumbents compare to entrants in their market. In Table H.1, I restrict the sample to treatment market-days from the entry experiment in which I observe take-up. I then run the following specification to compare entrants to incumbents in their market-day for various outcomes $Y$:

$$Y_{idw} = \alpha + \beta E_{idw} + \boldsymbol{\lambda_{dw}} + \epsilon_{idw} \tag{29}$$

where $Y_{idw}$ is the outcome for trader $i$ in market $d$ in week $w$, $E_{idw}$ is a dummy for whether trader $i$ is an entrant, and $\boldsymbol{\lambda_w}$ are market-day fixed effects. Standard errors are clustered at the trader-level (the source of variation in $E_{idw}$).

Table H.1 presents the comparison. I do not see any statistically significant differences in terms of quantity sold or price at which sold between the entrants and incumbents. Entrants appear to offer the same product as incumbents, with no statistical differences in credit provisions or quality. Sensibly, entrants are much less likely to report knowing other traders in the market well (ranked on a scale of "well," "somewhat well," and "not well"). Of particular interest, given the overall finding that entrants appear to be colluding with incumbents, entrants are no less likely to report discussing the price or agreeing on a price with other traders.

Finally, Table H.2 presents the IV effect of the number of traders on prices. Column 1 presents the main specification, which includes all traders in control and treatment markets for the entry experiment. The column presents a point estimate of about a 1% drop in prices in response to one additional trader (however, again, note that this is not significant). Column 2 excludes the entrants, and therefore isolates the effect on just incumbents in entry markets. While there is no statistically significant difference between these two point estimates, I do observe a smaller point estimate in Column 2, at a little more than half that of Column 1. One should be cautious about over-interpreting results that are not statistically significant; however, it is some suggestive evidence that, of the limited price effect that is observed, part of this effect may be driven by entrants undercutting incumbents, while part may be from incumbents being driven to lower their prices in response to entry.

Table H.1: **Comparison between Entrants and Incumbents.** Point estimates on a dummy for being an entrant (compared to incumbents). The sample is restricted to market-days in which entry occurred.

| | **Point Estimate** | **SE** | **Baseline Value** | **Obs** |
|---|---|---|---|---|
| Sell anything | -0.06 | 0.07 | 0.88 | 481 |
| Ln Kgs | -0.14 | 0.27 | 5.62 | 425 |
| Ln Price | -0.01 | 0.01 | 3.37 | 412 |
| Quality (1-4,4=best) | 0.06 | 0.09 | 2.60 | 479 |
| % Credit | 0.00 | 0.02 | 0.03 | 430 |
| Know others well | -0.40 | 0.08 | 0.49 | 417 |
| Discuss price | 0.03 | 0.08 | 0.33 | 417 |
| Agree price | -0.00 | 0.07 | 0.26 | 417 |

Table H.2: **Effect of Entry on Incumbents-Only.** Column 1 presents the main IV specification from Table 6, while Column 2 presents the same specification with entrant traders removed from the sample.

| | (1)<br>Ln Price | (2)<br>Ln Price |
|---|---|---|
| Num Traders | -0.00955<br>(0.00582) | -0.00536<br>(0.00570) |
| Mean Dep Var | 3.364 | 3.366 |
| N | 1776 | 1691 |
| Sample | All | Incum. only |
| Market FE | Yes | Yes |
| Week FE | Yes | Yes |

# I Appendix: Entry Offer Take-up and Simulated Effects

Table I.1 presents the first stage effect of the entry offer on the number of traders and the reduced form effect of the entry offer on log price broken down by the number of traders in each market.[56]

I generally observe higher levels of entry in larger markets, though this trend is non-monotonic. Note that the point estimate on markets with ten traders seems to suggest an increase in the number of traders that is outside the known bound of three, the total number of entrants given the entry offer. There is only one market in this bucket, and therefore this discrepancy is likely due to noise not absorbed by the market and week fixed effects. The other estimates, however, appear to be in a reasonable range.

[56]The number of traders is defined as the average number of traders observed in the market over the course of the experiment. In order to remove any increases in the number of traders driven by the entry experiment, this figure uses the average of the predicted number of traders each week, based on market and week fixed effects.

Table I.1: **Take-up of Entry Offers by Market Size.** The outcome variable is regressed on week fixed effects, market fixed effects, dummies for the number of traders present in the market (as categorized at baseline in Figure A.2), and interactions of each of these dummies and an indicator for entry treatment. Only interaction term coefficients are displayed here.

| | (1)<br>Num Traders | (2)<br>Ln Price |
|---|---|---|
| 1 Trader | -0.0213 | -0.0229 |
| | (0.209) | (0.0166) |
| 2 Traders | 0.283 | 0.00102 |
| | (0.123) | (0.00933) |
| 3 Traders | 0.295 | -0.00822 |
| | (0.160) | (0.00768) |
| 4 Traders | 0.796 | -0.00628 |
| | (0.325) | (0.00711) |
| 5 Traders | 1.083 | -0.0278 |
| | (0.338) | (0.0126) |
| 6 Traders | 0.910 | -0.00941 |
| | (0.570) | (0.00423) |
| 7 Traders | 0.205 | 0.0191 |
| | (0.209) | (0.00718) |
| 8 Traders | 1.760 | -0.0535 |
| | (0.119) | (0.00499) |
| 9 Traders | 1.080 | -0.00273 |
| | (0.321) | (0.00597) |
| 10 Traders | 8.425 | -0.0394 |
| | (0.142) | (0.00521) |
| Type | FS | RF |
| Mean Dep Var | 4.305 | 3.364 |
| N | 2045 | 1776 |
| Market FE | Yes | Yes |
| Week FE | Yes | Yes |
| Num Traders Control | Yes | Yes |

## J Appendix: Shift in Kernel Density of Sigma

Using 1,000 bootstrapped samples of my entire data, I calculate for each the baseline competitiveness parameter $\sigma_C$ (as determined by the estimated pass-through rate in that sample and the estimate of demand curvature $\delta$ from the corresponding bootstrapped demand parameter estimates). I then estimate, for each sample, an estimate of $\sigma_T$ under entry, using each sample's estimate of $\sigma_C$, the price effect of entry, and $\delta$ and $a$ from the corresponding bootstrapped demand parameter estimates. The kernel densities of $\sigma_C$ and $\sigma_T$ are plotted here. A Kolmogorov-Smirnov test cannot reject that these two distributions are the same (D=0.0183, p-val=0.996).

Figure J.1: **Kernel density of sigma control and sigma treatment**.