American Economic Review 2020, 110(12): 3705–3747
https://doi.org/10.1257/aer.20171397

Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya†

By Lauren Falcao Bergquist and Michael Dinerstein*

African agricultural markets are characterized by low farmer revenues and high consumer food prices. 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. Randomized cost shocks and demand subsidies are used to identify a structural model of market competition. Estimates reveal that traders act consistently with joint profit maximization and earn median markups of 39 percent. Exogenously induced firm entry has negligible effects on prices, and low take-up of subsidized entry offers implies large fixed costs. We estimate that traders capture 82 percent of total surplus. (*JEL* L13, O13, Q11, Q12, Q13)

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, Barrett, and Minten 2009; Burke, Bergquist, and Miguel 2019). High

* Bergquist: Department of Economics, University of Michigan (email: lfb@umich.edu); Dinerstein: Kenneth C. Griffin Department of Economics, University of Chicago (email: mdinerstein@uchicago.edu). Esther Duflo was the coeditor for this article. We thank Hoyt Bleakley, Marshall Burke, Ying Fan, Benjamin Faber, Willa Friedman, Benjamin Handel, Kaushik Krishnan, Nicolas Li, Jeremy Magruder, Jamie McCasland, Edward Miguel, Craig McIntosh, Odyssia Ng, Katalin Springel, Christopher Sullivan, Michael Walker, and Dean Yang, as well as seminar audiences at Boston University, Columbia, Cornell Dyson, IFPRI, Microsoft Research, MIT/Harvard, Northwestern, NYU, Penn State, UC Berkeley, UC Santa Cruz, University of Chicago, University of Michigan, University of Tennessee, University of Wisconsin, World Bank Development Research Group, and Yale, and to participants at many workshops and conferences. We also thank three anonymous referees for their valuable comments. Finally, we thank Honglin Li, Meshack Odhiambo Okello, Ben Wekesa, Deborrah Muthoki Wambua, and Innovations for Poverty Action for excellent research assistance. Bergquist would like to acknowledge financial support from the NSF Graduate Research Fellowship and the research initiative "Private Enterprise Development in Low-Income Countries" (EG 1764 and 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 AEARCTR-0000682 (see Bergquist 2015). IRB approval provided by the University of California, Berkeley (CPHS Protocol Number 2014-08-6612) and the Maseno University Ethics Review Committee (Proposal Reference Number MSU/DRPC/MUERC/000081/14). The data repository is openicpsr-119743 (see Bergquist and Dinerstein 2020). All errors are our own.

† Go to https://doi.org/10.1257/aer.20171397 to visit the article page for additional materials and author disclosure statements.

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 form of competition among intermediary traders in agricultural markets in developing countries. Whether, and how much, 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, such as 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 a high degree of market power, gains from policies that reduce traders' operating costs may be captured mostly 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, we present some of the first experimental evidence on the market structure in which African agricultural traders operate. To this end, we implement three randomized control trials that deliver new empirical evidence on the extent of cost pass-through, the shape of demand, and the effects of entry on market prices. We interpret these estimates in the context of a structural model and estimate the model of imperfect competition that best fits the data and its welfare implications.

In the first experiment, we exogenously reduce traders' marginal costs by offering to all traders in a market a substantial, month-long subsidy per kg sold. We then observe how much of this reduction in costs is passed through to the price offered to consumers. We find that traders pass through only 22 percent of this reduction in costs to customers.

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, we run a second experiment to estimate the curvature of demand. In this experiment, we offer consumers random reductions in price spanning a range of counterfactual pass-through rates and measure the resulting quantities purchased. This experimental variation in prices allows us to identify consumer demand without having to rely on strong identification assumptions. We specify and estimate a highly flexible demand function.

To determine agricultural intermediaries' form of competition, we start with a simple model of demand and supply that nests Cournot competition and joint profit maximization. This model transparently maps the experimentally estimated pass-through rate and demand curvature into inference about the form of competition. We find that traders are not competing: the estimated parameter governing how traders value other traders' profits 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. We can rule out Cournot competition with 95 percent confidence.

We then relax key assumptions in the simple model and estimate a more general model that allows for within-market trader heterogeneity and non-constant marginal costs, as well as an extensive margin to demand. We rely on the same experimental variation to identify the more general model, and introduce additional instruments describing how a trader's choices in one market depend on experimental variation in the other markets in which he sells. These multi-market traders, and their exposure to experimental variation across all of their markets, allow us to identify a non-constant marginal cost function. Our estimates from the more general model again imply conduct consistent with joint profit maximization; as with the simple model, we can reject Cournot competition. Traders capture high markups, with the median trader earning a 39 percent markup. We also offer non-nested tests of joint profit maximization and Cournot competition that again point toward joint profit maximization as describing traders' conduct.

Our third experiment generates exogenous entry by offering traders incentives to enter randomly selected markets for the first time. These results serve two purposes. First, we test whether policies that encourage market entry can decrease market power and promote competition. Second, traders' willingness to accept the entry subsidies, the size of which is randomized, reveals how they trade off fixed costs versus variable profits, which enables us to estimate total trader profits and conduct a complete welfare analysis. We find that the entry subsidies induce an additional 0.6 traders per market-day on average, a 16 percent increase over the mean market size (and 21 percent over the median). These additional traders have only small impacts on price. We find that competition in markets in which entrants have no prior connections increases to a level that is statistically indistinguishable from Cournot competition, but that markets in which entrants have prior connections continue to be collusive. Entrants with no prior connections, however, are less willing to take up the entry offer, indicating that the most likely compliers from a policy aimed at increasing entry may not be effective in increasing competition.

We use this experimental variation to identify an entry model in which potential entrants decide to enter a market when their variable profits upon entry, which depend on marginal costs and the impact of entry on market competition, exceed their fixed costs, net of the experimental entry subsidy. Our estimates indicate that potential entrants have high fixed costs that are positively correlated with marginal costs. Extending the estimates to incumbent traders, we find that the median trader's fixed costs constitute 71 percent of variable profits; the median (mean) trader keeps 12 percent (25 percent) of revenues as profit, but the largest traders earn the highest markups, such that in the aggregate traders capture 82 percent of total surplus, while consumers capture just 18 percent.

We use our estimated demand, quantity-setting, and entry models to solve for counterfactual equilibria were traders to engage in Cournot competition. We estimate that switching from joint profit maximization to Cournot competition would have large effects on surplus division, as consumer surplus would triple and deadweight loss would fall by one-third.

This paper is one of the first to experimentally test the competition model of rural agricultural markets directly. Previous attempts to measure competition or market efficiency 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, Barrett, 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 by 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 we study, in which over 95 percent 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 model of competition. These have generally found that average trader profits are moderately large, though subject to significant variability, leaving a question mark on whether these returns represent rents or risk premia (Dillon and Dambro 2016). Moreover, these direct measures are subject to potentially severe measurement error in the face of difficult-to-quantify search, own labor, and fixed costs, as well as a general lack of record keeping (Fafchamps, Gabre-Madhin, and Minten 2005).[2] The sensitivity of directly asking about profits in an environment in which traders are often labeled as exploitative presents a further challenge.

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 costs or whether farmers are simply unable to use this information to improve their bargaining position. There is therefore a paucity of causal evidence on how traders compete (Dillon and Dambro 2016) despite a growing interest in the role these intermediaries play in determining the allocation of gains from trade (Antràs and Costinot 2011; Bardhan, Mookherjee, and Tsumagari 2013).[4]

More broadly, this paper follows a long literature on testing among different models of competition. We follow Nevo (2001) and Miller and Weinberg (2017), among many others, in testing among common models of competition nested in a single framework. We also implement non-nested tests of joint profit maximization and Cournot competition in the spirit of Berry and Haile (2014) and similar to

[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 form 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 ours) is crucial to interpreting the pass-through rate, we experimentally estimate this curvature.

[2] Only 58 percent of traders in our sample keep any written records and, among this group, most records are fairly rudimentary.

[3] The exception is Hildebrandt et al. (2015), which finds that farmers who receive price information earn 5 percent 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.

Backus, Conlon, and Sinkinson (2019a), which avoids some of the issues with the nested approach. To identify the form of competition, we rely on an experimental cost shock. This idea of identifying and quantifying deviations from perfect competition using pass-through rates dates back to Sumner (1981), who used cost variation from cigarette excise taxes. Sullivan (1985) and Ashenfelter and Sullivan (1987) also used excise tax pass-through, while the trade literature has leveraged exchange rate pass-through (Feenstra 1989, Knetter 1989, Aw 1993, Goldberg 1995, Goldberg and Knetter 1999).[5] Bulow and Pfleiderer (1983) clarified the role of demand functional form in inferring competition from pass-through, and more recent contributions to the literature on pass-through include Weyl and Fabinger (2013) and Atkin and Donaldson (2015). Unlike many of these papers, we use experimental variation to identify the model components, adding to the recent literature using experimentally estimated parameters to understand market structure in developing countries (Keniston 2011).

This paper proceeds as follows. Section I describes maize markets in Kenya. Section II introduces a general and simplified theoretical supply and demand model. These motivate the experimental design, which we describe in greater detail in Section III. Section IV presents results on pass-through and demand, and then uses these estimates to identify the form of competition within the simple model. Section V relaxes some of the simple model's assumptions and presents results from the more general model. Section VI describes results of the entry experiment and estimation of an entry model that we use to quantify welfare in Section VII. Section VIII concludes.

I. 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 one-third of average gross caloric intake. The median household spends 9 percent of its annual expenditure on maize (and the poorest decile spends 14 percent). On the production side, about one-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.

Online Appendix 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 percent of their maize from small-medium farmers (selling < 5 tons), 16 percent from large farmers (selling ≥ 5 tons), and 33 percent from other traders. They buy primarily from counties throughout western Kenya and neighboring regions in eastern Uganda (the latter is more common during the lean season, when local supply is scarce in Kenya).[6] 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 repackaging, and then carry it onward to their destination of sale. In

[5] Relatedly, Panzar and Rosse (1987) developed predictions based on pass-through of cost shocks into revenues.

[6] Due to differences in crop calendars, farmers in eastern Uganda harvest maize several months earlier than those in western Kenya.

our sample, 64 percent of sales take place in open-air markets in rural communities. Sixteen percent is sold to millers, who grind maize into flour for sale to stores that serve urban consumers. Another 16 percent is sold to other traders, who sell in other areas of Kenya or eastern Uganda. A small portion of sales, about 2 percent, is sold to restaurants, schools, and other institutions. Finally, 2 percent 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.

A. *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 the maize trade.[7] However, engaging in large-scale, regional wholesale trade still requires significant working capital in order to pay for inventory, storage facilities,[8] and transport vehicles.[9] 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 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.

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

B. *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

[7] 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 US$100/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 US$2 "cess" tax charged to traders in the market each day.

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

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

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. Online Appendix 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 96 percent of traders reporting working in that market most weeks and only 1 percent saying that this was their first time in the market (see Online Appendix Table A.1). Seventy-seven percent have worked previously with *all* other traders in the market that day. As a result, 68 percent say they know the other traders in the market that day "very well," 26 percent "somewhat well," and only 6 percent "not very well." When asked directly, only 38 percent of traders report "discussing a good price" with other traders and only 30 percent report engaging in an explicit price agreement with other traders; the vast majority claim they are rigorously competing on price. However, 72 percent of traders work in a market in which at least one trader has reported the existence of a price agreement that day. Most (83 percent) traders visit only one of our 60 sample markets during the 12-week experimental period, though multi-market traders visit an average of 2.7 markets.

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. We therefore model consumers as considering one local market.[10] The median customer buys maize for consumption every week; storage is rare (see online Appendix C.1). The product itself is fairly homogeneous (see online Appendix C.2).

II. Theoretical Framework

The experimental design employed in this study is tightly tied to theory. In this section, we introduce a model of supply and demand for a homogeneous product, starting with a general model with limited assumptions. We then offer a simplified version of this model which demonstrates that, with the addition of a few key assumptions, a parameter nesting models of competition is a function of a small number of sufficient statistics. These sufficient statistics form the basis for the experimental design, with each of the three experiments implemented here designed to identify a specific parameter. Experiment 1 identifies pass-through, while Experiment 2 identifies the curvature of demand. We use these two estimates to infer the model of competition that best describes how traders operate. In Section V, we relax these simplifying assumptions and return to estimating the more general model, still relying on the experimental variation for identification. We then return to the third experiment, where the number of traders in the market is experimentally manipulated, and estimate the effect on how traders compete. Figure 1 provides an overview of how the experiments map to theory.

[10] Consistent with this, we find that the nearest market is on average 8km away (7km for the median market). On the same day, the nearest active market is 22km away (17km for the median market). Data on consumer behavior are drawn from a phone survey with 165 consumers randomly selected from the demand experiment sample. This survey was conducted in July and August 2016 immediately following data collection for the main experiment. In online Appendix H we find some evidence of substitution from markets outside of our sample into markets in our sample and discuss how this might affect our results.

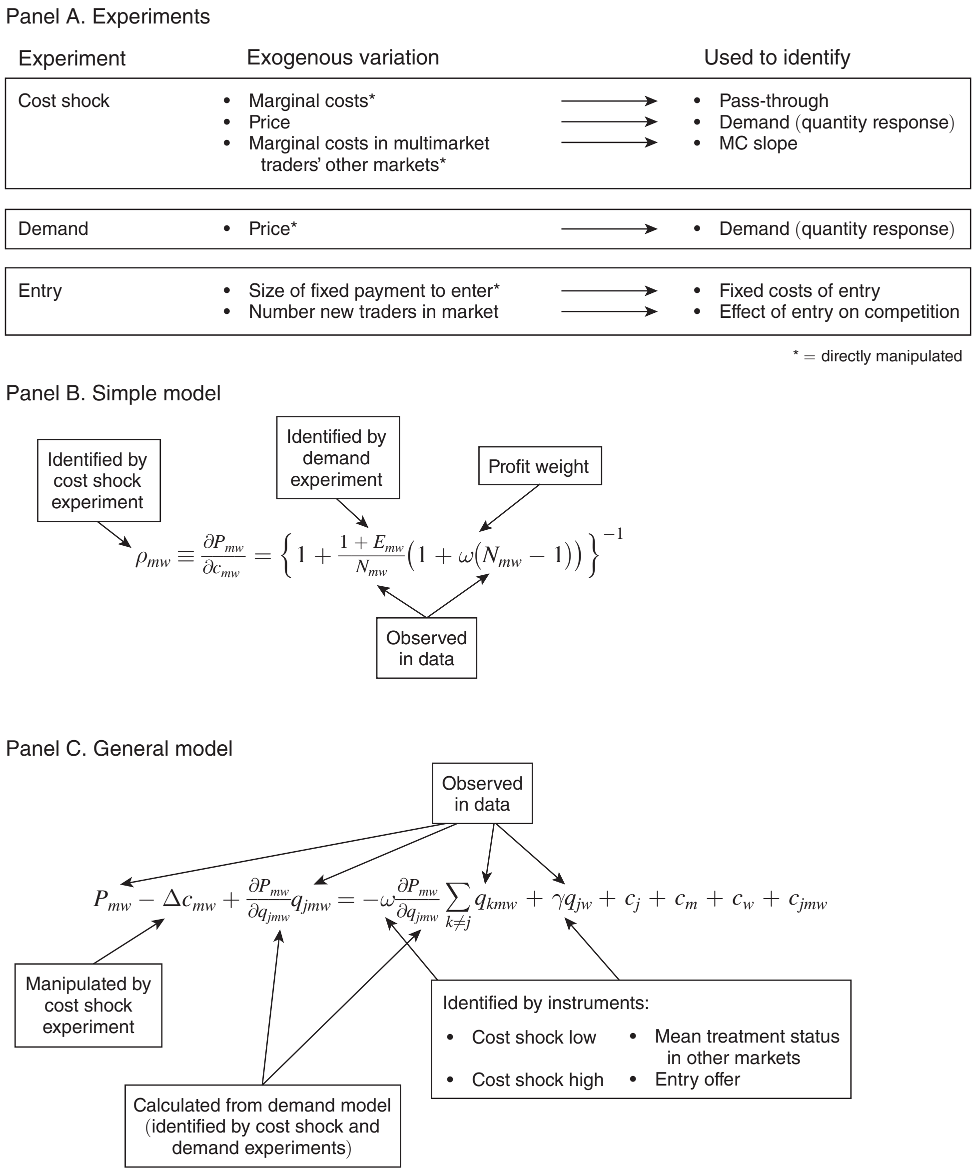

Figure 1. Experimental Design

Notes: Panel A lays out the experiments, the exogenous variation driven by each experiment, and the objects identified by said exogenous variation. Panel B and panel C lay out the main identifying equations for the simple and general model, respectively, and how each component is identified.

A. *Model Setup*

We begin with a model of demand and supply of a homogeneous good.[11]

[11] Whether the product is differentiated matters for how we infer the form of competition from equilibrium prices and quantities, as the same equilibrium responses to shocks could be consistent with one model of competition under homogeneous products and a different model of competition under differentiated products. However, the

Household i, in week w and its local market m, demands $q_{imw}(P_{mw}) \geq 0$ kilograms of maize at market price P_{mw}. Given the set $\mathcal{I}_{mw}$ households in local market m in week w, market demand is $Q_{mw}(P_{mw}) = \sum_{i \in \mathcal{I}_{mw}} q_{imw}(P_{mw})$ with corresponding inverse market demand $P_{mw}(Q_{mw})$. Traders choose quantities for each market they visit each week to maximize weekly profits.[12] We write trader j's maximization problem as[13]

$$(1) \qquad \max_{\{q_{jmw}\}_{m \in \mathcal{M}_{jw}}} \pi_{jmw} = \sum_{m \in \mathcal{M}_{jw}} P_{mw}(q_{jmw})\, q_{jmw} - C_{jw}(q_{j1w}, \ldots, q_{jMw}) + \omega \left[\sum_{m \in \mathcal{M}_{jw}} \sum_{k \neq j} P_{mw}(q_{jmw})\, q_{kmw} - \sum_{k \neq j} C_{kw}(q_{j1w}, \ldots, q_{jMw}) \right],$$

where $\mathcal{M}_{jw}$ is the set of markets j visits in week w and C_{jw} is j's total costs in week w across all markets. The second set of terms in brackets, multiplied by ω, captures other traders' profits in the same markets. Therefore $\omega \in [0,1]$ serves as a profit weight summarizing how a trader internalizes the profits of other traders in his market, with $\omega = 0$ reducing to Cournot and $\omega = 1$ reflecting joint profit maximization.[14] Taking the derivative of equation (1) with respect to the trader's quantity q_{jmw} in market m in week w yields the trader's first-order condition:

$$(2) \qquad P_{mw} = \frac{\partial C_{jw}}{\partial q_{jmw}} - \frac{\partial P_{mw}}{\partial q_{jmw}} \left(q_{jmw} + \omega \sum_{k \neq j} q_{kmw} \right),$$

where $\omega = 0$ and $\omega = 1$ reduce this to the familiar Cournot and monopoly first order conditions, respectively.

We will estimate a parametric form of this general model in Section V. But to highlight how our experiments can directly identify the model of competition, we first offer a complementary simplified version of the model that reduces inference

weight of evidence suggests maize sold in these markets is fairly homogeneous. There is little variation in quality, price, or other dimensions of the sale (e.g., credit is rarely used). See online Appendix C.2 for further detail. Price discrimination also appears quite limited, with an intra-cluster correlation of 0.9 between the prices that a trader offers his various customers throughout the day (see online Appendix C.3). This is likely because negotiations between traders and consumers are conducted in public, thereby limiting traders' ability to engage in dramatic price discrimination. We therefore assume a single market price.

[12] Specifying a weekly cost function allows for the possibility that the subset of traders active in multiple in-sample markets in the same week may have cost interdependencies across markets.

[13] The model employed here is static. While maize is in theory storable, empirically, consumer stockpiling is quite limited (see online Appendix Section C.1). Related work in the region suggests that credit constraints limit households' ability to arbitrage these price fluctuations (Burke, Bergquist, and Miguel 2019).

[14] Under this formulation, ω can be interpreted structurally, with traders directly valuing the profits of other traders, or as a parameter summarizing conduct. For the structural interpretation, the literature has focused on cases where there is a common claimant on agents' profits (e.g., vertically integrated units in Crawford et al. 2018 or firms with common investors in Backus, Conlon, and Sinkinson 2019b). In our setting, there are several features that could motivate a structural interpretation of the profit weight. Specifically, the structural profit weight encompasses relevant cases including traders from the same extended family, altruism toward fellow traders, and social insurance. If we proceed without a structural interpretation, the profit weight has no direct interpretation outside of $\omega = 0$ and $\omega = 1$. Fan and Sullivan (2018) derive a profit weight model consistent with a set of supergames, though they point out an additional term capturing rivals' deviation profits is necessary. We therefore primarily employ ω as a convenient formulation for nesting the two well-defined model of competition for which we will test empirically in our experiment. We also provide a direct non-nested test of Cournot and joint profit maximization. Finally, while versions of the reduced form approach can bias estimates toward too much competitiveness if the traders are colluding at a price below the monopoly price (Corts 1999), our empirical estimates in Section V indicate joint profit maximization.

about the model of competition to a few statistics. This relies on two simplifying assumptions. First, consistent with Fafchamps, Gabre-Madhin, and Minten (2005), we assume marginal costs ($c_{jmw} = \partial C_{jw}/\partial q_{jmw}$) 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 online Appendix Section D). In the general model, we will relax this assumption (though estimates do suggest costs are fairly flat). Second, we assume symmetry across traders in a market-week, specifically with respect to marginal cost ($c_{jmw} = c_{mw}$). Again, we will relax this assumption in the general model.

Given these assumptions, our first order condition reduces to

$$P_{mw} = c_{mw} - \big(1 + \omega(N_{mw} - 1)\big)\frac{\partial P_{mw}}{\partial Q_{mw}}\frac{Q_{mw}}{N_{mw}}, \tag{3}$$

where N_{mw} is the number of traders in the market m and week w. The term ω continues to nest Cournot competition and full collusion.[15] We see that equilibrium markups depend on the shape of demand and two features of market structure and trader behavior: the number of traders N_{mw} and whether these traders interact according to Cournot competition ($\omega = 0$) or joint profit maximization ($\omega = 1$).

B. *Pass-Through and Demand Curvature*

In the first part of the paper, we use two experiments, one identifying pass-through and one identifying consumer demand, to estimate ω in the observed market equilibria.[16]

Our first experiment estimates pass-through by experimentally reducing traders' marginal costs. To link this to the model, we take the derivative of equation (3) with respect to c_{mw}, which yields

$$\rho_{mw} \equiv \frac{\partial P_{mw}}{\partial c_{mw}} = \left\{1 + \frac{1 + E_{mw}}{N_{mw}}\big(1 + \omega(N_{mw} - 1)\big)\right\}^{-1}, \tag{4}$$

where $E_{mw} \equiv \big(Q_{mw}/(\partial P_{mw}/\partial Q_{mw})\big)\big(\partial(\partial P_{mw}/\partial Q_{mw})/\partial Q_{mw}\big)$ is the elasticity of the slope of inverse demand. Under the specific models of competition tested here, this equation reduces to

$$\rho_{mw} = \begin{cases} \left\{1 + \frac{1 + E_{mw}}{N_{mw}}\right\}^{-1} & \text{when Cournot competitive} \\ \{2 + E_{mw}\}^{-1} & \text{when collusive.} \end{cases} \tag{5}$$

[15] This equation also nests the solution to the Bertrand price-setting model, with $\omega = -1/(N_{mw} - 1)$. We will not focus on this case, but point out that because Bertrand competition implies a negative ω, rejecting Cournot with a positive point estimate also implies rejecting Bertrand.

[16] We will use experimental variation across markets to estimate ω. This relies on the additional assumption that all sample markets operate under the same form of competition. In addition to being necessary for statistical power, this assumption also reflects the case in which multi-market contact leads traders to make potentially coordinated decisions at a level above the market. We are also assuming that the treatment itself, which will be a cost shock, does not change ω. The shock is unlikely to introduce any unfamiliar traders to the market. We investigate this further in online Appendix Section I, where we find a very small, marginally significant effect that is likely to bias us toward Cournot competition. We relax this assumption when studying entry and allow ω to vary with whether traders have connections to each other.

Therefore, the level of pass-through ρ_{mw} depends on the competitive structure of markets ($\omega = 0$ or $\omega = 1$) and the curvature of demand E_{mw} (as well as on the number of traders N_{mw} in the Cournot model).[17]

Accordingly, our first experiment identifies pass-through and our second experiment identifies demand curvature. We then insert these experimentally estimated parameters into equation (4) from the sufficient statistics model and back out the profit weight ω in these markets. In Section V, we will estimate ω in the more general model that relaxes the assumptions of constant marginal cost and symmetric traders. Finally, we provide a direct non-nested test of Cournot and joint profit maximization that does not depend on specifying a continuous parameter that nests Cournot competition and joint profit maximization.

C. *The Effect of Entry on Competition*

A commonly prescribed policy intervention to increase market competition is to encourage entry by new market participants. In the second part of this paper, we evaluate the impact of such policies and estimate whether the model of competition responds to changes in the number or identity of market participants. We then use these entry decisions to estimate a structural model of entry into a market. These estimates let us quantify the distribution of trader fixed costs and allow us to conduct counterfactual exercises where trader participation responds endogenously to the form of competition.

III. Experimental Design

A. *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 online Appendix Section E for additional details on the sample selection procedure.

The two market-level experiments (cost shock and entry) were each run for four weeks in a row. This time spans about one-quarter of the full selling season in the region (March to July). This duration of treatment was selected to represent a relatively 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 2). 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 (cost shock treatment, entry treatment, and control) in a random order to allow for the inclusion of these fixed effects. Online Appendix Figure E.1 shows the six possible orders.[18] Each four-week block was broken by a one-week break during which the demand experiment was run in a subset of markets.

[17] Under Bertrand competition, $\rho_{mw} = 1$

[18] This randomization was first blocked by the day of the week of the market (done primarily for logistical ease as the trader cost-shock and entry treatment required additional management time to facilitate payments, and equal

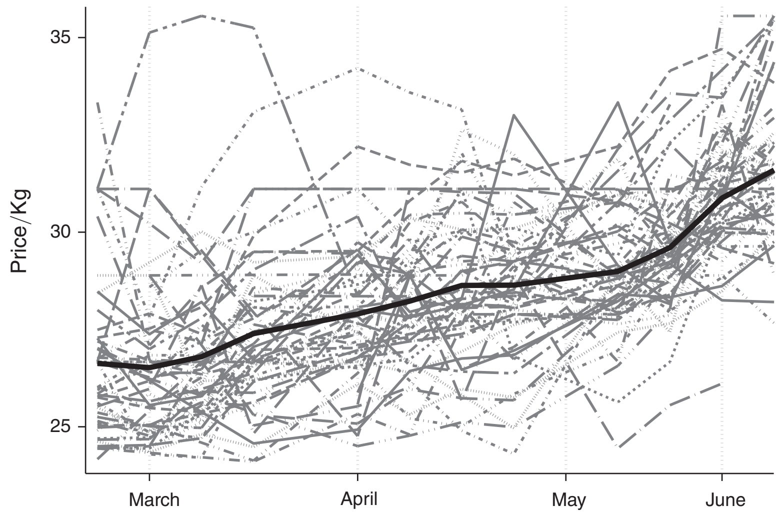

Figure 2. Maize Prices in Study Markets

Notes: Gray lines show the price for each market over the 12-week study period. The black line shows the average price across markets.

B. *Experiment 1: Trader Cost Shock Experiment*

In treatment market-days for the trader cost shock 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 5:00PM (after the market conclusion). Maize sold during the enumerators' presence in the market was eligible for the subsidy.[19] 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

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 online Appendix Section E for further details on this census.

[19] 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 percent 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 percent of traders.

who were absent in the previous week. All traders in the market therefore received an identical reduction in their marginal costs.

The 60 markets in the sample were divided into two groups: 45 markets received a "low" subsidy level of 200Ksh/90kg bag when they were in the cost shock treatment and 15 markets received a "high" subsidy level of 400Ksh/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 percent of the average price, while the "high" subsidy represents 15 percent of the average price. Payments were made via mobile money twice a day. See online Appendix Section J.1 for additional details about this experiment.

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.

C. *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.[20] 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 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 negotiation 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 percent of the cost-reduction subsidy had been passed-through in the cost shock experiment. Per 90kg bag, they were: {0, 25, 50, 100, 150, 200, 250, 300, 350, 400} Ksh. See online Appendix Section J.2 for additional details about this experiment.

[20] 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 sufficient statistics model is that these customers would have exhibited the same curvature of demand as the customers observed in the sample. For the general empirical model, we directly model the extensive margin and identify substitution on the extensive margin using results from the trader cost shock experiment.

D. *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 (i) to increase the probability that at least one trader took up the offer and (ii) 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, we then identified the set of eligible traders who (i) had never before worked in that market and (ii) did not work in other study markets that occur on the same day of the week in order to avoid inducing exit in our sample. The median market had 37 eligible traders, the minimum had 28, and the maximum had 56. From each of these sets, we then randomly selected the three traders who would receive the entry offers.

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 (US$49), a "medium" offer of 10,000Ksh (US$99), and a "high" offer of 15,000Ksh (US$148). The trader was eligible to receive this amount each time he visited the entry market on any of four offer days.[21] Payout was contingent on a few factors, of which traders were made aware during the offer call. They were that the trader must (i) come to the specified market on the specified date; (ii) arrive with a truck and at least 15 bags; (iii) stay for at least one hour and show intention to attempt sales. Payment was made via mobile money 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, a short follow-up phone survey was conducted, in which information was collected about the trader's activities on the day of the offer regardless of whether he accepted the offer. See online Appendix Section J.3 for additional details about this experiment.

E. *Data*

Data were collected in an identical way in all markets and in all periods (cost shock treatment, entry treatment, and control). Depending on the activity level of

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

each market, enumerators were assigned to monitor 1–4 traders.[22] 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 nontraditional 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.[23] Maize quality data were also collected for each trader (see online Appendix Section C.2 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.[24] 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.

IV. Estimating the Form of Competition in Simple Empirical Model

In this section we estimate the elements, or sufficient statistics, that enter equation (4) in the simplified model. We start by estimating pass-through and then turn to the curvature of demand. We end by backing out the implied model of competition.

A. *Pass-Through*

To measure pass-through, we estimate

$$P_{jmw} = \beta CC_{mw} + \gamma_w + \zeta_m + \epsilon_{jmw}, \tag{6}$$

where P_{jmw} is trader j's quantity-weighted average price in market m and week w,[25] CC_{mw} is the level of cost change per kg offered in market m on week w (i.e., CC is the *negative* value of the marginal cost subsidy in cost shock treatment markets and zero elsewhere, such that $CC_{mw} = \{0 \text{ Ksh}, -200 \text{ Ksh}, -400 \text{ Ksh}\}$), and γ_w and ζ_m are week and market fixed effects, respectively, included to improve

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

[23] These nontraditional reductions in price were uncommon, but they do add 1–2 percentage points to our estimate 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 discreet method of deviating from price agreements maintained with fellow traders. It may also help traders to make price changes more continuous (as prices otherwise typically change in increments of 50 Ksh/bag or 5 Ksh/goro, a 2.2kg tin).

[24] 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.

[25] Because the ICC of price within a trader in a given market-day is high (0.9), in practice there is little variation in the prices entering into this average.

Table 1—Pass-Through

	Price (1)	Price (2)
Cost change	0.224	
	(0.0434)	
Cost change–low		0.219
		(0.0538)
Cost change–high		0.228
		(0.0618)
Mean dependent variable	28.92	28.92
Observations	1,860	1,860
Market fixed effects	Yes	Yes
Week fixed effects	Yes	Yes

Notes: Price regressed on *Cost change*, the level of cost reduction per kg offered in market m on week w. *Cost change* is the negative value of the marginal cost subsidy in cost shock treatment markets and zero elsewhere, such that *Cost change* $= \{0$ Ksh, -200 Ksh, -400 Ksh$\}$) or $\{0$ US\$, -1.98 US\$, -3.96 US\$$\}$. Week and market fixed effects are included to improve precision. The sample includes traders in market-days in which the market was in either the cost shock treatment or control period: market days assigned to the entry treatment are omitted. Under this specification, the coefficient on the cost reduction term yields the pass-through rate, or $\partial P/\partial c$. The first column shows the overall pass-through rate of 22 percent. The second column shows pass-through rates separately by "low" and "high" offers.

precision. The sample includes traders in market-days in which the market was in either the cost shock treatment or control period: market days assigned to the entry treatment are omitted. Under this specification, the coefficient of interest is β, which yields the pass-through rate, or $\partial P/\partial c$. Throughout we cluster standard errors by market-block, the unit of randomization, and weight the regressions by the inverse number of traders in the market to give equal weight to each market.

To measure heterogeneity in the pass-through rate by the level of the cost change, we estimate

$$P_{jmw} = \beta_1 CC_{mw} \times Low_{mw} + \beta_2 CC_{mw} \times High_{mw} + \gamma_w + \zeta_m + \epsilon_{jmw}, \tag{7}$$

in which Low_{mw} ($High_{mw}$) is a dummy indicating whether the market was in a low (high) subsidy market. This allows for nonlinearities in the effect of the subsidy per kg. For other measures of heterogeneity, we run specifications similar to equation (7), conditioning on the desired dimension of heterogeneity.

Table 1 presents the main results of the pass-through experiment. In column 1, we see that pass-through is 22.4 percent, significantly different from zero at the 1 percent 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 help inform the functional form assumptions in the following section on demand estimation.

We explore heterogeneity by the number of traders in the market.[26] Figure 3 presents these results, which show little evidence of meaningful heterogeneity.

[26] This is the main source of heterogeneity prespecified in a design registry submitted prior to the beginning of the experiment. The number of traders is defined as the average number of traders observed in the market over the

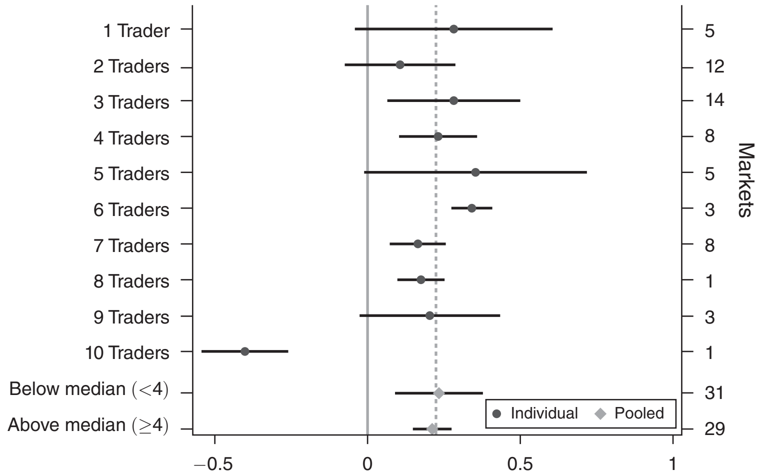

FIGURE 3. PASS-THROUGH BY MARKET SIZE

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

Estimates of pass-through rates are fairly tightly centered around the overall estimate of 22 percent 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. That pass-through does not vary with number of traders is consistent with, though not definitive evidence of, collusion (see equation (5)).

We further explore in Figure 4 other dimensions of heterogeneity that could matter for how traders compete.[27] First, we measure whether pass-through is different for markets on and off paved (tarmac) roads, which serves as a proxy for market geographic isolation. We find no evidence of heterogeneity by this measure. Next, we explore whether a higher intensity of explicit collusion predicts lower pass-through rates, measured by the number of market-days within a market in which traders have explicitly admitted to collusion.[28] The point estimates suggest that pass-through is similar across these markets, and the differences are not statistically significant.

B. *Demand Estimation*

As described in Section II, in order to draw inference in our simplified model about the level of competition from the observed pass-through, one must first understand the curvature of demand. To do so, we use the demand experiment to estimate how

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.

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

[28] We 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. We then divide the sample into markets above and below the median of this measure.

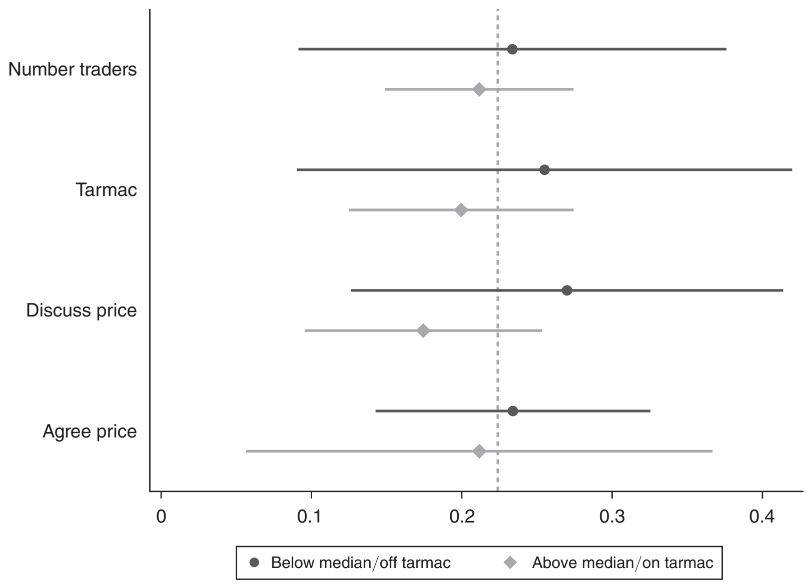

Figure 4. Pass-Through by Various Factors

Notes: Pass-through as estimated in markets in each category (bars represent the 95 percent confidence interval). The average for the full sample is 22 percent (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.

a household's quantity of maize demanded varies with price. We plot the estimated treatment effects by randomized price change in Figure 5. Consumers were given the option to choose a new quantity to purchase at the new price. In the left graph, we plot the treatment effects on the fraction of consumers that changed quantities once offered the discount. This fraction is increasing in the size of the price change and reaches nearly 30 percent at the highest discount point. These changes translate into a strong relationship between price and quantity changes, as shown in the right graph.

Many common demand functional forms impose curvature and thus would divorce inference about how traders compete from the data. Instead, we will estimate demand curvature. Ideally we would estimate demand nonparametrically, but because household demand is so heterogeneous and demand curvature depends on a function of the second derivative of demand, inference based on nonparametric demand estimation is highly imprecise. We thus impose a flexible class of demand functions that nests several commonly used functional forms, while leveraging the panel structure of the experiment to maximize statistical power in the presence of demand heterogeneity. We use the design feature that customers were approached after agreeing on an initial price and quantity combination, so that we observe customers making two choices with an exogenous price difference. Let t index the "experimental period," i.e., whether the customer has yet to be approached $(t = 0)$ or already been offered a subsidy level $(t = 1)$. This allows us to capture the considerable household demand heterogeneity while still estimating demand

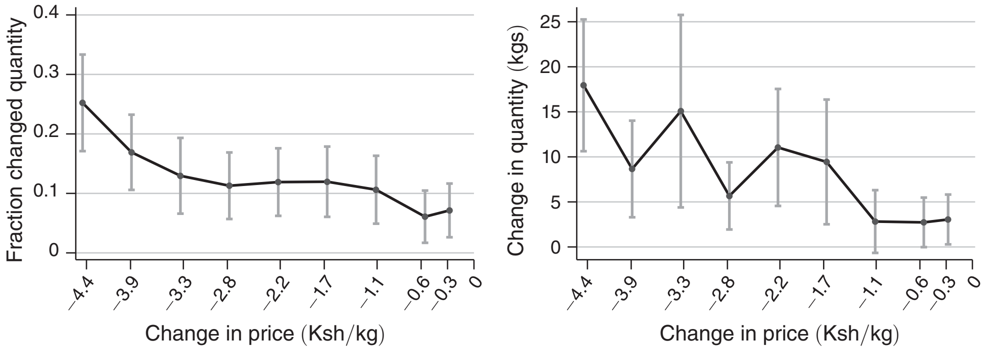

Figure 5. Demand Experiment Quantity Effects

Notes: The panels show treatment effects on the fraction of consumers who choose a new quantity (left) and change in quantity transacted (right) by level of price change. Quantity changes are expressed in levels and are means that include consumers making no change. Price changes (Ksh/90kg bag) are determined by the randomized subsidy, which can take up to 10 values: $\{0, -25, -50, -100, -150, -200, -250, -300, -350, -400\}$ Ksh/90kg bag. The x-axis on the plots are price changes per kg. The mean initial quantity and price are 65kg and 32 Ksh/kg, respectively. The 95 percent confidence interval is shown around each point estimate.

parameters precisely. We embed this household heterogeneity within a general Bulow-Pfleiderer class of demand functions. Household i's demand q_{imt} is

$$q_{imt}(P_{imt}) = \begin{cases} \left(\frac{a - P_{imt}}{b_i}\right)^{\frac{1}{\delta}} \eta_{imt} & \text{if } P_{imt} \leq a \\ 0 & \text{if } P_{imt} > a, \end{cases} \tag{8}$$

where $a, b_i, \delta, \eta_{imt} > 0$.

The variable a is the choke price, such that households are indifferent between purchasing or not; b_i captures persistent (within the few minutes the experiment lasts) household heterogeneity while η_{imt} is a high-frequency demand shock. Here the price faced (P_{imt}) varies by customer, even within a market, because the experiment varies the size of the discount. Given individual demand $q_{imt}(P_{mt})$, market demand is $Q_{mt}(P_{mt}) = \sum_{i \in \mathcal{I}_{mw}} q_{imt}(P_{mt})$.

We choose this particular class of demand functions for its flexibility, tractability, and empirical foundation. First, this demand structure is flexible, nesting many of the functional forms common to the development and trade literature, including linear demand and quadratic demand.

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).[29] Recall that in the simple model described in Section II, the pass-through rate is determined by the profit weight ω, potentially the number of traders, and the slope of inverse demand E. While in theory E can vary with quantity, this is

[29] This property holds both for this specification of individual and market demand. Atkin and Donaldson (2015) also rely on this property in estimating markups.

a second-order term for which it is already difficult to get precision on a single estimate using the full pooled data (as we will show below); attempting to further estimate E at different quantity levels 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 market demand function is

$$P_{mt} = a - \kappa_{mt} Q_{mt}^{\delta}, \tag{9}$$

where constant κ_{mt} depends on the household heterogeneity (b_i) and time shocks (η_{imt}) in market m. In this case, the elasticity of the slope of inverse demand, E_{mw}, reduces to $\delta - 1$. Therefore, equation (4) simplifies to

$$\rho_{mw} \equiv \frac{\partial P_{mw}}{\partial c_{mw}} = \left\{1 + \frac{\delta}{N_{mw}}\big(1 + \omega(N_{mw} - 1)\big)\right\}^{-1}, \tag{10}$$

with our specific models of competition reducing to

$$\rho_{mw} = \begin{cases} \dfrac{N_{mw}}{N_{mw} + \delta} & \text{when Cournot competitive} \\ \dfrac{1}{1 + \delta} & \text{when collusive.} \end{cases} \tag{11}$$

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 (10), because E is constant across q, this class of demand functions predicts a constant pass-through rate for a given ω, 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, we are able to test for this prediction of constant pass-through. Because markets are randomized into receiving the low versus high subsidy rate, 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.

Our focus on demand functions with E constant in q motivates our specification of consumer heterogeneity. We place heterogeneity in the b_i term while keeping a the same for all consumers.[30] This preserves constant E, matching the high versus low cost shock pass-through results and providing us with the benefits of tractability when estimating a higher-order term, but comes at the expense of modeling an endogenous extensive margin.[31] Given the potential importance of extensive margin

[30] We also impose constant δ across consumers. Because δ relates to a higher-order term in demand, we lack the variation to estimate heterogeneous δ with any precision.

[31] Variation in a, the choke price, would drive variation in which consumers enter the market at different prices. Constant a implies either full or no endogenous participation. The number of consumers may still vary for exogenous reasons, but given our functional form, would not affect inference on the model of competition.

responses to price, in Section V we add heterogeneity in a and discuss the functional form we impose.[32]

Estimation, Identification, and Results.—To estimate the demand model, we implement a log transformation of equation (8) and take the first difference within customer:

$$\log(q_{im1}) - \log(q_{im0}) = \frac{1}{\delta}\big(\log(a - P_{im1}) - \log(a - P_{im0})\big) + \big(\log(\eta_{im1}) - \log(\eta_{im0})\big), \tag{12}$$

where $P_{im1} - P_{im0}$ is the subsidy amount.

Any endogeneity in prices due to persistent demand conditions is accounted for by the panel nature of this specification. We further use the randomized reduction in the price paid by consumers from the demand experiment as an instrument for the price term $(\log(a - P_{im1}) - \log(a - P_{im0}))$ to address any remaining endogeneity (e.g., in high-frequency demand shocks, η_{imt}) as we identify the model's parameters.

We run the analysis with 936 observations. We estimate the vector of parameters $\Theta = (a, \delta)'$ in equation (12) using generalized methods of moments. We construct our IV sample moments as the vector of 9 dummy variables for each positive discount level times the residual.[33] We minimize the GMM objective function using the optimal weighting matrix from two-step estimation.

Results are presented at the top of Table 2, which show the point estimates and 95 percent confidence intervals. Note that the confidence interval on δ is wide (for example, we cannot rule out very curved inverse demand of $\delta = 6.43$). This is because δ, which represents the elasticity of the slope of inverse demand (plus 1), is a higher order object which we are underpowered to measure with great precision, even with over 900 observations from the demand experiment. That said, we have enough precision to rule out some standard demand functional forms, including linear demand ($\delta = 1$). Moreover, this degree of precision is sufficient for our purposes. As we will see in the next subsection, from the point estimate on δ, we can predict the level of pass-through that one should expect under various models of competition; we will find the prediction of one model to line up closely with what is observed empirically. Further, even at the bounds of our estimate of δ, we can still reject that what we see empirically is consistent with other common models of competition.

[32] Another approach to integrating an extensive margin would have been to model market choice and quantity demand jointly in a discrete-continuous demand system, where the discrete and continuous choices are linked via Roy's Identity. While point estimates from this model lead to similar conclusions about the model of competition, the standard errors are large. Furthermore, given our data, summarizing the discrete choice of where to buy requires considerable dimension reduction, such as through a logit functional form assumption. This assumption puts strong restrictions on demand curvature that we seek to avoid.

[33] We drop the consumers offered no discount as they have no within-consumer price variation and do not make a new quantity choice.

Table 2—Model Estimates

	Parameter estimate	95% confidence interval lower bound	95% confidence interval upper bound
Simple model			
a	42.76	41.56	43.96
δ	4.07	1.71	6.43
ω	0.78	0.05	7.48
General model			
μ_a	29.15	28.84	52.19
σ_a	2.87	2.48	60.00
δ	4.21	0.70	9.64
ω	1.07	0.20	3.09
γ	0.0006	−0.0006	0.0016
General model, $\omega = 1$			
γ	0.0007	−0.0010	0.0016

Notes: The point estimates and the 95 percent confidence intervals for the estimated parameters of the simple and general models are displayed. The final row is the point estimate for the marginal cost slope under the general model where we impose $\omega = 1$. This is the relevant cost parameter for markup and welfare analysis, which we implement given that traders maximize joint profits. The 95 percent confidence interval on γ in the model with $\omega = 1$ imposed does not account for the testing step. For a specification that allows for joint estimation of γ and testing of the model of competition, see online Appendix Section B.

C. *Model of Competition*

First, we 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 δ of 4.1, we use equation (10) to estimate the average pass-through rate one should expect to observe in the experiment under various models of competition. If markets are Cournot competitive ($\omega = 0$), we should observe pass-through rates that vary with the number of traders: $\rho = N/(N+4.1)$.[34] Given the distribution of number of traders in each market in our sample, the expected pass-through rate if markets are Cournot competitive is 46 percent. On the other hand, if markets are collusive ($\omega = 1$), we should expect to observe 20 percent pass-through.[35]

Figure 6 displays the bootstrapped distribution of ρ.[36] We see that the mass of the distribution of ρ is concentrated near the predicted pass-through of 20 percent under collusion. The dotted lines, which identify the 95 percent confidence interval, clearly reject a ρ consistent with that predicted under a model of Cournot competition (and Bertrand competition).

This exercise does not take into account the fact that δ is estimated imprecisely. To account for this imprecision, we generate a bootstrapped distribution of ω by

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

[35] While not nested in our model, if traders are Bertrand competitors, we should observe 100 percent pass-through, at least in this simple model.

[36] The distribution was constructed using 1,000 block bootstrapped samples where blocks are defined by market $\times$ 4-week-blocks (the level at which treatment was randomized). There are 180 such clusters from 60 markets.

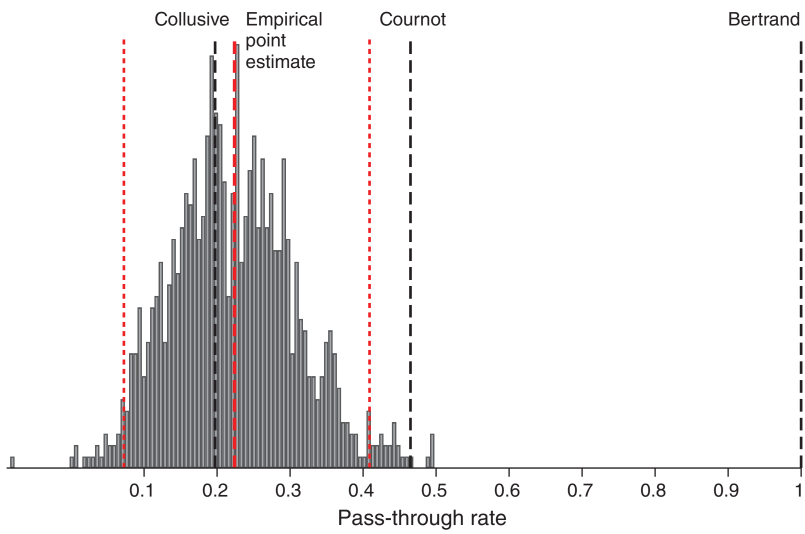

FIGURE 6. PREDICTED PASS-THROUGH UNDER THREE SIMPLE MODELS

Notes: Given our demand curvature estimate, we predict within our simple model that one would have observed 100 percent pass-through in a Bertrand competitive market, 46 percent pass-through in a Cournot competitive market, and 20 percent pass-through in a collusive market environment. The distribution of empirical pass-through, calculated for 1,000 bootstrapped samples, is shown in gray. The point estimate and 95 percent confidence interval are shown in red.

estimating equation (10) with 1,000 bootstrapped estimates of ρ and δ.[37] The top graph in Figure 7 presents this distribution, overlaid with the benchmark values of ω under Cournot competition and collusion. We plot in red the value of ω predicted by the point estimates on ρ and δ. The point estimate of ω is 0.78, with a 95 percent confidence interval of $(0.05, 7.48)$, which is quite close to, and statistically indistinguishable from, the model benchmark of $\omega = 1$ under perfect collusion. Moreover, while the collusive market benchmark of $\omega = 1$ lies well within the 95 percent confidence interval, the levels of ω predicted by a Cournot model and perfectly competitive model lie outside these bounds. We are therefore able to reject them with 95 percent confidence.[38]

Under this simple model, the observed pass-through rate is therefore consistent with an underlying market structure in which traders act as if they maximize joint profits.

V. Estimating the Form of Competition in the General Empirical Model

In Section II we started with a general model before imposing two supply-side assumptions: trader symmetry and constant marginal costs. We also made a

[37] We take the expectation of equation (10) over our markets with heterogeneous numbers of traders and estimate ω to make equation (10) hold in expectation.

[38] We are able to reject a Cournot competitive model because the confidence interval around δ, however large, does exclude the extreme curvature necessary to justify such low pass-through under a Cournot model. To achieve a predicted ρ of 22 percent under a Cournot model, we would have required a δ of about 12.

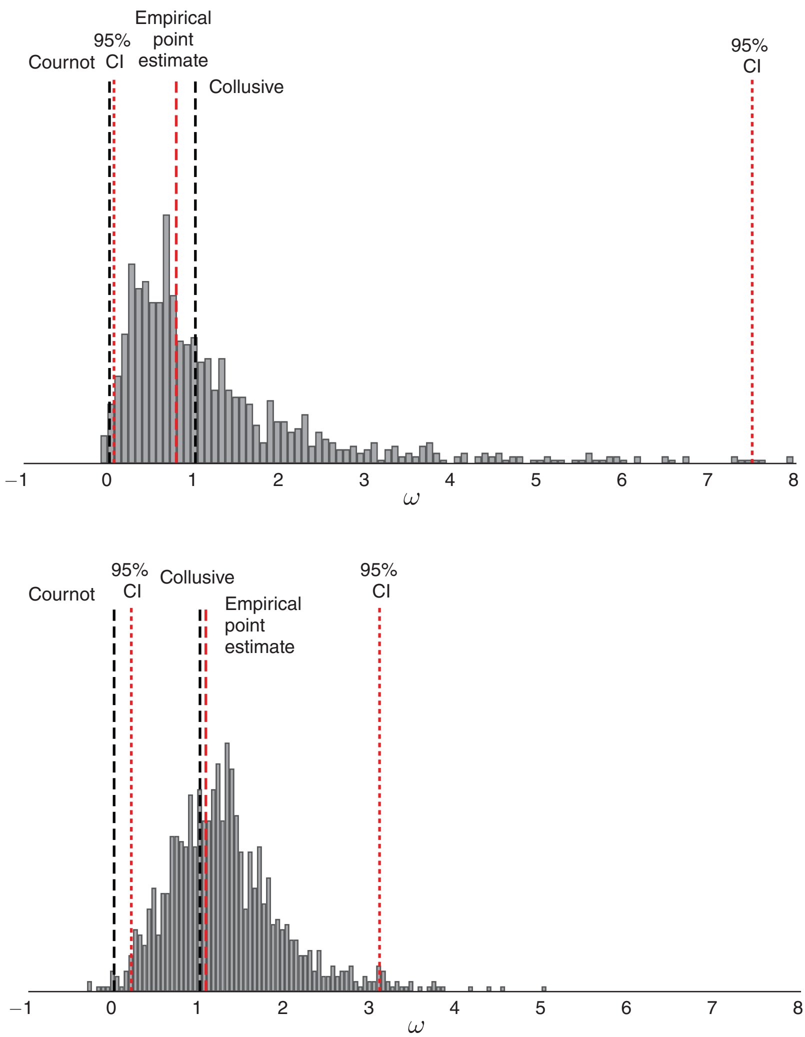

Figure 7. Profit Weight (Omega) Estimates: Simple and General Models

Notes: The top panel shows the bootstrapped estimated ω distribution for the simple model while the bottom panel is for the general model. Bootstrap estimates of ω come from 1,000 bootstrapped estimates of ρ and δ plugged into equation (10) for the simple model and full demand and supply bootstrap for the general model. Recall that $\omega = 0$ if Cournot competitive and $\omega = 1$ if perfectly collusive. The point estimate and 95 percent confidence interval are shown in red.

simplifying assumption in the demand model, placing heterogeneity across consumers in b_i, rather than in a, which restricts adjustment on the extensive margin. The value of these simplifying assumptions is that they allow us to reduce each experiment to identifying a single parameter; this one-to-one matching of experiment to theory makes transparent the link between experimental results and the identification of the model of competition. Using this approach, Section IV estimated the profit weight and found evidence consistent with joint profit maximization.

To test whether our conclusions are robust to these simplifying assumptions, we now relax the supply-side assumptions and extend our demand model to accommodate further household heterogeneity and an extensive margin. We view this more general model as complementary as it relaxes assumptions but with a less direct link between experimental results and inference on the model of competition, albeit still using experimental variation for identification of all parameters.

A. *General Demand Model*

We start with the general demand model as estimated demand will be an input into our supply model. We continue with the model in Section IVB but add further consumer heterogeneity by letting a, the choke price, be consumer-specific:[39]

$$q_{imt}(P_{imt}) = \begin{cases} \left(\frac{a_i - P_{imt}}{b_i}\right)^{\frac{1}{\delta}} \eta_{imt} & \text{if } P_{imt} \leq a_i \\ 0 & \text{if } P_{imt} > a_i, \end{cases} \tag{13}$$

where $a_i, b_i, \delta, \eta_{imt} > 0$.

This heterogeneity makes the model more flexible and introduces a heterogeneous household extensive margin. If the price is above a household's a_i, then the household will not purchase any maize at the market that week. This heterogeneity, though, proves challenging to identify even with the panel structure of the demand experiment. We thus impose a functional form on the heterogeneity, assuming $a_i \sim N(\mu_a, \sigma_a^2)$. This heterogeneity permits more flexible market demand changes for a price change.[40]

Further, because the demand experiment reaches consumers only after they chose to visit the market and agree to a price-quantity bundle, we supplement our demand experiment results with additional data moments from the cost shock experiment to bring in more information about consumers' choices related to the extensive margin.[41] Before specifying the moments, we analyze the effects of the cost shock on market-level transacted quantities. We estimate reduced-form regressions, using market-week level versions of equations (6) and (7) but replacing the dependent variable of price with four quantity measures: number of transactions, transaction rate,[42] kgs per transaction, and total kgs. We summarize the results in Table 3.

We find a large quantity response to the cost shock. For a decrease in marginal cost of 1 Ksh/kg, we estimate an increase of 846 kgs in a market/day (column 5 of Table 3). The effect is driven almost entirely by the extensive margin, as we see an increase of 11 transactions (column 1). The effect on the intensive margin, which combines increased demand from a fixed set of consumers and compositional

[39] To the extent that some consumers may make multiple transactions, we allow a to vary by transaction.

[40] While individual demand remains in the Bulow-Pfleiderer class of demand functions, market demand is no longer in the class and we will no longer rely on the sufficient statistics formula.

[41] The variation in the demand experiment alone is technically sufficient to identify μ_a and σ_a, the parameters of heterogeneity in a, as a_i affects not just whether to buy anything but also how much; however, because the additional moments better target extensive margin choices, we present them as our main specification. Note, though, that we get very similar results when relying only on the demand experiment moments.

[42] We observe the number of transactions for each market-day. We convert these to a transaction rate by dividing by the maximum daily number of transactions observed in sample for each market. Estimates of the profit weight in Section VB are robust to increasing the market size by at least a factor of 2.

Table 3—Quantity Effects

	Number of transactions (1)	Number of transactions (2)	Transaction rate (3)	Kgs/ transaction (4)	Kgs (5)
Cost change	−11.05		−0.0871	3.143	−845.6
	(1.9083)		(0.00931)	(4.111)	(163.5)
Low treatment		33.43			
		(5.522)			
High treatment		36.49			
		(10.94)			
Mean dependent variable	60.92	60.92	0.603	87.79	4,809.5
Elasticity	−19.0		−16.0	6.0	−20.4
Observations	454	454	454	454	454
Market fixed effects	Yes	Yes	Yes	Yes	Yes
Week fixed effects	Yes	Yes	Yes	Yes	Yes

Notes: The table presents reduced-form results. *Number of transactions* is the number of transactions observed in any week for a given market. *Transaction rate* is the number of transactions divided by the maximum number of transactions observed in any week for a given market. *Kgs/transaction* is the average kgs per transaction in any week for a given market. *Kgs* is the total kgs sold in any week for a given market. *Cost change* is the change in cost (Ksh/kg) from treatment: −2.22 in the low cost shock treatment and -4.44 in the high cost treatment. *Low treatment* and *High treatment* are dummy variables for whether the market-week is in the low cost shock and high cost shock treatments, respectively. Elasticities are calculated using estimated demand curves in online Appendix Table H.1 and evaluated at the mean price and mean quantity under the low cost shock (the median of the three experimental groups).

changes in the set of consumers transacting, is much smaller and not statistically significant. To convert these estimates to elasticities, we estimate

$$Q_{mw} = \beta P_{mw} + \gamma_w + \zeta_m + \epsilon_{mw} \tag{14}$$

and instrument for market price with dummy variables for being in the low cost shock treatment or the high cost shock treatment. These two-stage least squares estimates, presented in online Appendix Table H.1, show steep demand curves. The implied elasticities when evaluated at the mean price and mean quantity under the low cost shock, the median of our three experimental groups, are listed at the bottom of the table. For total quantity (*Kgs*) we estimate an elasticity of −20.4, while we estimate an extensive margin (*Num Trans*) elasticity of −19.0.

It is worth noting that these estimates do not necessarily imply large consumption elasticities. The object estimated here is a market residual demand elasticity, the percentage increase in quantity that a market's traders can expect to sell for each percentage decrease in price, and is the relevant elasticity for characterizing traders' strategic incentives. This may be distinct from the increase in quantity that a consumer would consume if she faced lower prices universally. We explore the potential explanations for the large market residual demand elasticity in online Appendix Section H.

Moreover, these simple specifications assume linear demand. However, column 2 of Table 3, which shows the reduced-form effect of each treatment on the number of transactions, reveals that doubling the cost reduction from low to high produces only a slightly greater number of transactions. In fact, if we calculate the elasticities piecewise, we find that the low cost shock versus control comparison yields an estimated elasticity of -28.1, while the low cost shock versus high cost shock comparison yields an elasticity of just -3.3. This suggests that demand may be highly nonlinear (or in our demand model, a large fraction of potential consumers have a_i near market prices). Unfortunately, the cost experiment alone has insufficient variation to estimate a nonlinear model; this motivates our combination of the cost shock experiment variation with the demand experiment's variation, with 9 exogenous price points, in estimating demand with a flexible functional form.

We thus add six moments from the cost experiment to our demand estimation: the mean market-week transaction rate and kgs per transaction in each experimental arm: control, low cost shock, high cost shock. We include the experiment's effect on the mean market-week transaction rate to capture extensive margin changes. And given that the consumers induced by the lower prices to purchase may be different from the consumers who show up even at high prices, we also include the moments related to kgs per transaction. We keep separate moments for the low and high cost shocks to capture the nonlinearities evident in Table 3. The sample moment values are in online Appendix Table F.1.

Using these moments and the high-powered demand experiment IV moments, we estimate demand using three-step GMM with an optimal weighting matrix and bootstrapped standard errors. We report the estimates in the middle of Table 2. We estimate δ to be 4.21, close to the estimate from our simpler demand model. For the distribution of a_i, we estimate a mean of 29.15 and a standard deviation of 2.87. This tight distribution is consistent with the large observed response in transaction rates to the cost shocks. The estimated demand model generates a mean demand elasticity, at market prices and quantities, of -3.2. These elasticity estimates are lower in magnitude than the implied elasticities from the linear model in Table 3, though they are consistent with the estimate from the low cost shock versus high cost shock comparison. The estimated demand model also yields declining kgs per transaction, driven by lower demand from marginal consumers, which matches the (noisy) results on kgs per transaction from the cost experiment variation.

B. *General Supply Model*

Turning to our supply model, we now relax the previous assumptions on trader symmetry and constant marginal costs. To do so, we estimate our general supply model from equation (2) with a functional form for total cost that allows marginal costs to vary on several dimensions, including with respect to quantities. We specify trader j's total costs in week w as

$$(15)\quad C_{jw} = \frac{1}{2}\gamma q_{jw}^2 + \sum_{m \in \mathcal{M}_{jw}} \Big(FC_{jmw} + \big(c_m + c_w + c_j + \Delta c_{mw} + c_{jmw}\big) q_{jmw}\Big),$$

where Δc_{mw} is the experimental cost shock, q_{jw} is j's total weekly quantity, and FC_{jmw} is j's fixed cost from trading in market m in week w. By modeling costs as depending quadratically on the total quantity sold in a week, we allow for increasing (or decreasing) marginal costs in the quantity sold in a given market-day, as well as the quantity sold throughout the week. Adding potential cost interdependencies across the week is important because many traders source maize on a weekly basis; we therefore allow for the possibility that marginal costs are increasing with quantities, due to, for example, the opportunity cost of not selling at another market that week or the cost of having to source from new farmers to fulfill demand at other markets that week. Our cost function also allows marginal costs to vary across markets (c_m) and weeks (c_w). This flexibility is important as certain markets may be more accessible or certain weeks may feature higher sourcing costs. Finally, we introduce heterogeneity across traders, allowing traders to face systematic differences in marginal costs (c_j). The remaining cost term, c_{jmw}, represents trader j's marginal cost shock in week w in market m (note this is distinct from the experimental cost shock, Δc_{mw}).

Plugging this cost function into equation (2) and taking the first-order condition with respect to quantities, we estimate the following supply model:

$$(16) \qquad P_{mw} - \Delta c_{mw} + \frac{\partial P_{mw}}{\partial q_{jmw}} q_{jmw} = -\omega \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} + \gamma q_{jw} + c_j + c_m + c_w + c_{jmw}.$$

We now have two primary parameters remaining to be estimated: ω, the profit weight equal to 0 under Cournot competition and 1 under joint profit maximization, and γ, the rate at which marginal costs increase (or decrease) with quantities.[43] We group components without unknown parameters on the left-hand side, noting that the right-hand side is now linear in the parameters ω and γ, which facilitates their estimation with linear model techniques. The cost shock c_{jmw} is our structural error.

Because other traders' quantities $(\sum_{k \neq j} q_{kmw})$ and own weekly quantity (q_{jw}) are determined in equilibrium, they are likely correlated with c_{jmw}. Therefore, we will rely on instruments to identify ω and γ. To maintain the close connection with the field experiments, we will construct our instruments using experimental variation only. Because treatment status was determined through experimental randomization, market-block treatments are orthogonal to trader type (specifically, c_{jmw}).

The first set of instruments are whether the market-week was randomized to have a low cost shock, a high cost shock, or no cost shock. While the cost shock (Δc_{mw}) is directly in the first-order condition we estimate, it has a known coefficient. We can thus rely on the (market-level) cost shock as an exogenous shifter of other traders' quantities as well as where the traders are on the demand curve $(\partial P_{mw} / \partial q_{jmw})$.

We choose the second set of instruments to target identification of non-constant marginal costs. Simultaneously identifying the model of competition and non-constant marginal cost is difficult and often depends on an instrument that rotates demand (Bresnahan 1982). We take a different strategy by looking for exogenous variation in a trader's total quantity sold in other markets that affects the focal market only through the cost function. In our setting, we exploit the fact that some traders operate in multiple in-sample markets and thus their exposure to cost

[43] We first estimate $\partial P_{mw} / \partial q_{jmw}$ using our demand estimates. Online Appendix Section F offers details on model estimation, including how we estimate $\partial P_{mw} / \partial q_{jmw}$ for each market-week.

shocks in *other* markets exogenously changes over time.[44] The experimental status of a multi-market trader j's *other* markets shifts trader j's total weekly quantity sold, which affects trader j's quantity supplied in the focal market through an increase (or decrease) in marginal cost, as determined by the size of γ. We thus construct the following instruments: the fraction of trader j's in-sample markets in week w that have a low cost shock or a high cost shock. To increase power and generate additional variation in quantities sold that week, we add more instruments to capture variation coming from the entry offer experiment. In that experiment, which we will discuss further in Section VI, potential entrants received low, medium, or high entry subsidies. The size of the subsidy predicts the likelihood of taking up the offer and entering a new market, which affects total quantity sold in a week. We thus use whether trader j randomly has a low, medium, or high entry offer (to another market) as additional instruments to identify non-constant marginal costs. We also add an instrument for the fraction of j's markets for which another trader has an entry offer.

Given these instruments, we run two-stage least squares to generate point estimates for ω and γ and construct confidence intervals by bootstrapping the entire demand and supply estimation. We present our estimates in the middle of Table 2 and the bootstrap distribution of ω in the bottom panel of Figure 7. We estimate ω to be 1.07, very close to the benchmark value of ω under joint profit maximization, with a 95 percent confidence interval of $(0.20, 3.09)$. Thus, we arrive at similar conclusions about the model of competition based on our more general model.

We can also directly test one of the assumptions of the sufficient statistics model: constant marginal cost. Our estimate of γ, the marginal cost slope, is 0.0006 Ksh/kg, and 0 is well within a fairly tight 95 percent confidence interval of $(-0.0006, 0.0016)$. This point estimate is small, implying that a 1 standard deviation increase in weekly (in-sample) quantity sold (2,300 kgs) corresponds to a cost increase of just 1.73 Ksh/kg. This pales in comparison to the heterogeneity in trader-market-week marginal cost intercepts (marginal cost for the first kg), where we estimate a standard deviation of 10.84 Ksh/kg.

In addition to the above approach, which tests whether ω coincides with well-defined models of competition, we can also implement a non-nested test of Cournot competition and joint profit maximization that does not treat ω as a structural parameter (Bresnahan 1987, Villas-Boas 2007). This test uses the exact same identifying variation as in the prior specification, but flips the approach. Rather than using the fact that the experimental treatment status is orthogonal to trader cost type (due to randomization) and then identifying the model of competition, it instead assumes a particular model of competition and then tests whether such independence holds, as it should under the correct model (Berry and Haile 2014, Backus, Conlon, and Sinkinson 2019a). Specifically, we modify equation (16), imposing either $\omega = 0$ (Cournot) or $\omega = 1$ (joint profit maximization) and inserting the cost shock Δc_{mw} on the right-hand side.[45] Under the correct model of competition,

[44] Though being a multi-market trader is unlikely to be random, by including trader fixed effects in our marginal cost specification, we isolate variation coming from the rotating experimental schedule over time.

[45] An added benefit is that this specification allows for joint estimation of γ and testing of the model of competition.

the randomized cost shock and trader cost type should be uncorrelated and thus the coefficient on the cost shock (π) should be zero:

$$P_{mw} - \Delta c_{mw} + \frac{\partial P_{mw}}{\partial q_{jmw}} q_{jmw} + \omega \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw} = \pi \Delta c_{mw} + \gamma q_{jw} + c_j + c_m + c_w + c_{jmw}. \tag{17}$$

Forming test statistics from the 1,000 bootstrap iterations, we reject $\pi = 0$ under Cournot competition, but fail to reject $\pi = 0$ under joint profit maximization. Results from this alternative approach are therefore consistent with our conclusions from the nested test. We provide more details in online Appendix Section B.

C. *Trader Markups*

We now turn to the implications of joint profit maximization for trader markups and variable profits.[46] Given within-market variation across traders in quantities transacted, our homogeneous goods model implies considerable cost, and thus markup, variation. We estimate a median markup of 39 percent and a mean of 48 percent, where we express markups as $(P_{mw} - MC_{jmw})/P_{mw}$ with MC_{jmw} being the marginal cost at the equilibrium quantity. There is wide dispersion in markups, with an estimated standard deviation of 35 percent.

Do these markups seem sensible? One check is to compare survey data on prices received by producers to those paid by final consumers. While we lack data on farmer prices in this study, a concurrent study conducted by Bergquist and McIntosh (2019) collected prices from local markets in neighboring eastern Uganda during the same four month period as this study. As much of the maize coming into western Kenya during the lean season comes from eastern Uganda,[47] prices observed in rural Uganda markets during this period approximate the relevant producer price.[48] Comparing this price with the mean consumer price observed in our data, we can generate a back-of-the-envelope markup estimate, which we find to be 45 percent, quite close to the median markup of 39 percent estimated from our model.

These large markups lead to large estimates of variable profits per market-day.[49] Variable profits are highly skewed with some traders earning very high variable profits. The median trader-market-week generates 3,400 Ksh while the mean generates 14,000 Ksh. These estimates indicate that traders charge prices well

[46] Note these estimates depend on the shape of demand and the cost function. Unless we interpret ω structurally, cost estimates only have meaning in the context of specific model. Given the evidence presented above, we therefore assume joint profit maximization ($\omega = 1$) and re-estimate the supply model to back out the parameters of the cost function. We show the new estimate of γ in the bottom row of Table 2. Unsurprisingly given our estimated ω is so close to 1, the estimate of γ is nearly unchanged at 0.0007 Ksh/kg.

[47] Uganda, known as the breadbasket of East Africa, experiences a harvest that occurs about four months earlier than that in western Kenya. Therefore, much of the maize available in our Kenyan study market during the four-month lean season in which this study is run is supplied by regions in eastern Uganda.

[48] While the markets studied in Bergquist and McIntosh (2019) are quite rural, these are not quite farmgate prices and therefore these prices are likely to be, if anything, slightly higher than the price eastern Ugandan farmers receive at farmgate. The estimates produced by this back-of-the-envelope exercise are therefore likely an underestimate of markups.

[49] For multi-market traders, we calculate weekly profits and divide by the number of in-sample markets to get a per-market estimate.

above their marginal costs. But whether traders actually make high total profits depends on fixed costs. We will use entry choices to learn about the magnitude of fixed costs in Section VI and thus defer our complete discussion of welfare until Section VII.

VI. Entry

Given that markets look fairly collusive, one natural policy response in the absence of antitrust enforcement is to encourage greater entry, especially among traders who might be willing to compete. There are several policies that could potentially encourage entry, such as offering lines of credit to potential new traders to rent long-haul trucks and disseminating information about profitable markets more broadly. But whether an entrant will introduce competition or collude is unclear. This is what we test in the third experiment, in which we randomly offer traders incentives to enter new markets.

Understanding entry is also relevant for evaluating the welfare implications of collusion. If low entry rates allow collusion to persist, then fixed costs of entry might be sufficiently large to dissipate high levels of variable profits. But because fixed costs do not affect the maximization of variable profits, the cost shock experiment tells us little about the distribution of fixed costs. Instead, we identify fixed costs by turning to the entry experiment that shifts how traders trade off fixed costs and variable profits.

A. *The Cost of Entry*

Because the offer amount is randomized, we can use traders' willingness to accept the offer as a measure of willingness to enter new markets. Table 4 presents take-up at each subsidy level (take-up defined as ever accepting any of the four market-day offers). Sensibly, we see that take-up increases in the size of the subsidy: take-up is 12 percent for the low offer (5,000 Ksh, or US\$49), 28 percent for the medium offer (10,000 Ksh, or US\$99), and 42 percent for the high offer (15,000 Ksh, or US\$148). The fact that take-up is far from universal, even with the high offer that exceeds mean variable profits, suggests that the cost of entry appears to be high in this setting.[50]

We first explore heterogeneity in willingness-to-enter by a few key variables prespecified 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, we estimate the following specification for the pool of 180 potential entrants:

$$(18) \qquad T_{jm} = \alpha + \beta X_{jm} + \epsilon_{jm}, \qquad \text{(i)}$$

$$T_{jm} = \alpha + \beta X_{jm} + \zeta_m + \epsilon_{jm}, \qquad \text{(ii)}$$

[50] 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, fewer than 5 percent of traders who did not take up the offer cite trust issues as the explanation.

Table 4—Take-Up of Entry Offers

	Offer Ksh	Amount US$	Take-up rate	Observations
Low offer	5,000	49	0.12	60
Medium offer	10,000	99	0.28	60
High offer	15,000	148	0.42	60

Notes: Offers ranged from 5,000–15,000 Ksh (US$49–148). *Take-up* = 1 if the trader *ever* took up an offer during any of the four weeks for which the offer was available.

in which T_{jm} is a indicator representing whether trader j ever took up an offer to enter his assigned market m; X_{jm} is the variable by which we explore heterogeneity. In specification (ii), we control for market fixed effects (ζ_m), such that we only look at differential take-up of the entry offer *within* the same market. We 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), we cluster standard errors by trader in both regressions.

Figure 8 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 the effect of distance on take-up to that of the offer amount, we estimate that an additional 50km in distance is roughly equivalent to a drop of US$46 in the offer amount.[51] Having contacts in the entry market is correlated with higher take-up (albeit not quite significantly). The point estimate suggests 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.[52]

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). Fifty-three percent of all markets had at least one day (out of four) with entry. Thirty-eight percent of all market-days had entry, 26 percent of which had more than one entrant. In total, an average entry market had an additional 0.6 traders present, an increase of 16 percent over the mean market size and 21 percent over the median.

[51] The precision of the distance effect drops 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.

[52] 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).

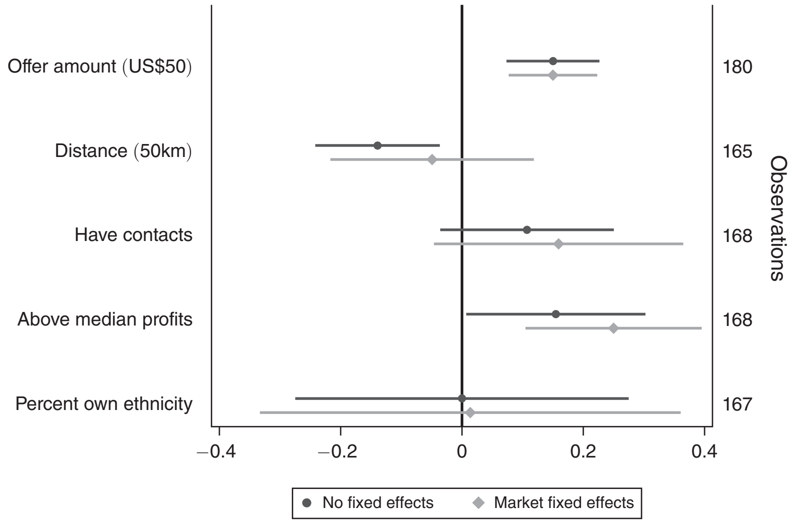

FIGURE 8. HETEROGENEITY IN WILLINGNESS-TO-ENTER

Notes: 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 percent confidence interval is plotted.

B. *The Effect of Entry on Price*

We turn now to the effect of entry on prices. To measure the reduced-form effect of the offer, we estimate

$$P_{jmw} = \alpha + \beta EntryMarket_{mw} + \gamma_w + \zeta_m + \epsilon_{jmw}, \tag{19}$$

where P_{jmw} is the average price per kg charged by trader j in market m in week w, $EntryMarket_{mw}$ is a dummy for whether market m is in an entry market in week w, and γ_w and ζ_m are week and market fixed effects, respectively. Standard errors are clustered at level of market $\times$ 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, cost shock treatment periods are omitted). Under this specification, the coefficient of interest is β, which yields the price reduction observed in the entry offer market.

We also run an IV specification to determine the effect of entry on prices:

$$P_{jmw} = \alpha + \beta NumEntrants_{mw} + \gamma_w + \zeta_m + \epsilon_{jmw} \tag{20}$$

in which $NumEntrants_{mw}$ represents the number of entrants in the market that day, for which we instrument with the $EntryMarket_{mw}$ dummy. Table 5 presents these results. We see a strong first-stage effect on the number of entrants (column 1), while the number of incumbents does not change (column 2). Reduced form effects are small and marginally significant, with only a 0.18 Ksh (or 0.6 percent of the

Table 5—Effect of Entry on Prices

	Number entrants (1)	Number incumbents (2)	Price (3)	Price (4)
Entry market	0.636	−0.0574	−0.180	
	(0.0601)	(0.131)	(0.106)	
Number entrants				−0.283
				(0.160)
F-statistic FS				111.9
Mean dependent variable	0.303	4.045	29.04	29.04
Observations	1,776	1,776	1,776	1,776
Market fixed effects	Yes	Yes	Yes	Yes
Week fixed effects	Yes	Yes	Yes	Yes

Notes: The variable *Entry market* is a dummy for treatment status in the entry experiment. *Number entrants* is the number of traders present in the market on that day who were offered a subsidy to enter. *Number incumbents* is the number of traders presented in the market on that day who were not offered a subsidy to enter. Column 1 presents the first-stage effect of treatment on the number of entrants. Column 2 presents the effect of the treatment on the number of incumbents. Column 3 presents the reduced-form effect of treatment on price. Column 4 presents the effect of the number of entrants on the price, instrumenting for the number of entrants with treatment.

mean) drop in prices (column 3). Column 4 presents the result of using treatment status as an instrument for the number of entrants. We see that the entry of one trader reduces prices by 0.28 Ksh (or 1.0 percent of the mean), with a p-value of 0.077.

C. *The Effect of Entry on Competition*

What does the small price decrease tell us about how the underlying competitive environment (summarized by ω) has changed? We have not directly modeled how ω is determined, but the introduction of new trader to a market could affect the stability of a collusive arrangement. We explore this using the entry offer experiment by estimating whether ω changes with entry.

Let ω_n be the profit weight if there is no entry and ω_e be the profit weight with entry of new traders. We return to our general supply model and let ω depend on entry:

$$(21)\qquad P_{mw} - \Delta c_{mw} + \frac{\partial P_{mw}}{\partial q_{jmw}} q_{jmw}$$

$$= -\omega_n(1 - Entry_{mw}) \frac{\partial P_{mw}}{\partial q_{jmw}} \sum_{k \neq j} q_{kmw}$$

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

where $Entry_{mw}$ is an indicator for whether market m has realized entry in week w.[53] We have already in prior sections estimated ω_n to be consistent with joint profit maximization, so we now impose $\omega_n = 1$ and leave ω_e to be estimated.

[53] We estimate profit weights that are constant within market-week. An alternative would be to let incumbent traders place different weights on other incumbents' profits and entrants' profits. We cannot separately identify these two models though our estimates lead to similar qualitative conclusions.

Realized entry likely depends on market and potential entrant characteristics. We thus instrument for entry with an indicator for whether market m in week w is part of the entry offer experiment. Because of the experimental randomization, this instrument is orthogonal to market characteristics. We restrict attention to incumbents (note that even if the market is randomized into treatment, the cost shocks, c_{jmw}, of the entrants themselves may still be correlated with treatment as entrants may have different costs from non-entrants or incumbents).[54] Our identification assumption then is that incumbents' cost shocks, which do not depend on market outcomes, are orthogonal to whether the market is an entry offer market this week.

We present our estimate of ω_e, with a bootstrapped confidence interval, in the top row of Table 6. We estimate ω_e to be 0.66 with a 95 percent confidence interval of $(0.10, 0.96)$. The point estimate falls in the middle of Cournot and joint profit maximization, and we can reject Cournot $(\omega = 0)$ and joint profit maximization $(\omega = 1)$. This indicates that the entry may change how traders compete.

We cannot formally distinguish whether the estimate of $\omega_e = 0.66$ reflects a yet-to-be-considered form of competition in all markets (e.g., collusion at a quantity above the perfect collusive quantity) or rather an average treatment effect arrived at from a subset of markets switching fully from joint profit maximization to Cournot competition.[55] However, we can generate suggestive evidence on this point by examining whether ω_e varies according to the prespecified sources of heterogeneity: number of preexisting contacts, trader size, and trader ethnicity.

While we find no significant heterogeneity based trader ethnicity and noisy heterogeneity based on trader size (see online Appendix Section G), we do find suggestive evidence of heterogeneity by whether the entrant has contacts in the market. Thirty-two percent of potential entrants have contacts in the entry market, and these potential entrants may find it easier to integrate into preexisting collusive arrangements. Along these lines, recall from Figure 9 that the offer take-up rate is higher for potential entrants with contacts than those without. We thus estimate whether entrants without contacts have differential effects, where we instrument for entrant type (whether he has contacts) with the type of the potential entrant that (randomly) received the high subsidy offer. We present the estimates in the last two rows of Table 6. We estimate that markets with entrants that have contacts have $\omega_e = 0.95$, quite close to the collusive benchmark of $\omega = 1$, with a 95 percent confidence interval of $(0.15, 1.53)$. Thus, entrants with contacts appear to integrate into the lack of competition already present in the markets. The story is different for entrants without contacts. Markets with these entrants have estimated $\omega_e = 0.44$, with a 95 percent confidence interval of $(-0.48, 0.96)$. For these markets, we can reject joint profit maximization but cannot reject Cournot competition.

These results suggest that entry's effectiveness in inducing more competition depends on the entrant type.[56] Policymakers interested in increasing competition may therefore prefer policy instruments that target entrants without contacts. However, this may be challenging to achieve, as the lower offer take-up rates for

[54] We still keep entrants' quantities as part of the sum of other traders' quantities sold.

[55] And unless we commit to a structural interpretation for ω, a value of 0.66 has no specific meaning.

[56] Whether a trader has contacts is not necessarily a fixed characteristic so it is possible that entrants that plan to visit a market well beyond a four-week period may find it worthwhile to form relationships that facilitate collusion.

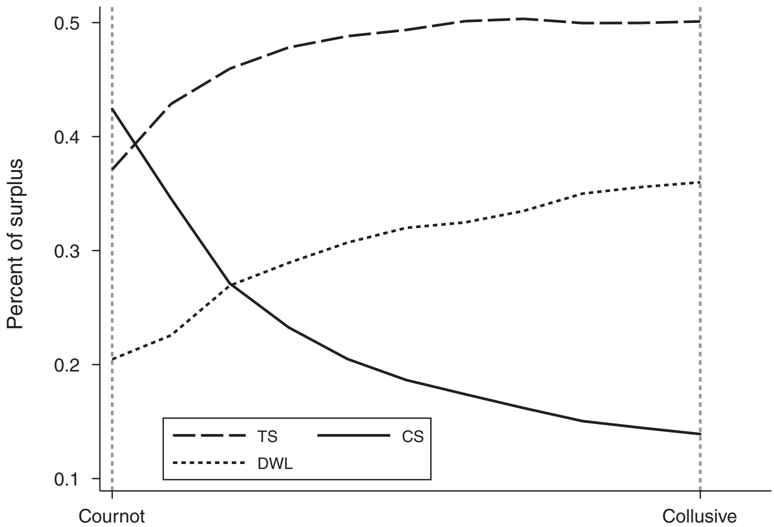

Figure 9. Welfare Counterfactuals

Notes: Counterfactual division of welfare is shown for the market-weeks in the cost shock experiment and the control group. The estimated division of surplus under joint profit maximization is shown at the far right vertical dotted line, suggesting that trader surplus (TS) is 50 percent of total surplus, while consumer surplus (CS) is only 14 percent and deadweight loss (DWL) is 36 percent. Movements to the left represent increases in competition. The dotted vertical line at *Cournot* indicates how this division would be altered if the market operated under Cournot competition.

Table 6—Effect of Entry on Competition

Group	Parameter estimate	95% confidence interval lower bound	95% confidence interval upper bound
Pooled model			
ω_e all entrants	0.66	0.10	0.96
Heterogeneous by contacts			
ω_e^{with} entrants with contacts (32 percent)	0.95	0.15	1.53
$\omega_e^{without}$ entrants without contacts (68 percent)	0.44	−0.48	0.96

Notes: The top row presents the estimate of the profit weight ω when a trader enters in the entry experiment. Markets that do not receive entry keep $\omega = 1$. The second and third columns show the bounds of the 95 percent confidence interval, calculated with 1,000 bootstrap iterations. The bottom two rows show separate profit weights depending on whether the entrant has contacts in the market. Of the potential entrants in the entry experiment, 32 percent have contacts in the targeted market.

entrants without contacts suggest these individuals are less willing to enter. This may be because, given that they induce more competition, these entrants likely earn lower profits than those able to facilitate join profit maximization.

But to draw any definitive conclusions on profit levels, of either incumbents or entrants, we must first estimate traders' fixed costs. While fixed costs are not directly observable in the data, they are key determinants of traders' entry decisions, which we do observe. We therefore estimate a model of traders' entry decisions to back out a fixed cost distribution and conduct welfare analysis. As with our earlier demand and supply models, we will focus on using exogenous variation to identify the model parameters, this time using the entry experiment.

D. *Entry Model*

We model potential entrant j as choosing to enter market m in week w if the variable profits from entering exceed the fixed cost (net of any experimental subsidy):

$$(22) \quad Entry_{jmw} = \begin{cases} 0 & \text{if } \pi^V_{jmw}\left(MC^0_{jmw}, \omega\right) < FC_{jmw} - EntrySubsidy_{jmw} \\ 1 & \text{if } \pi^V_{jmw}\left(MC^0_{jmw}, \omega\right) \geq FC_{jmw} - EntrySubsidy_{jmw}, \end{cases}$$

where π^V_{jmw} and FC_{jmw} are, respectively, trader j's variable profits and fixed costs from entering market m in week w, and $EntrySubsidy_{jmw}$ is the randomized subsidy from the entry experiment. We include MC^0_{jmw}, the marginal cost intercept $(c_j + c_m + c_w + c_{jmw})$, and ω as arguments to variable profits to highlight that entry decisions vary across traders based on heterogeneous fixed costs, heterogeneous marginal cost functions, randomized subsidy levels, and whether the trader changes the market's model of competition. For potential entrants with contacts in market m, we assume $\omega = 1$, and for potential entrants without connections, we assume $\omega = 0$.

The model clarifies the link between the entry experiment subsidy variation and observed entry outcomes. But it also demonstrates why the average treatment effect of an entry offer on entry decisions, or market outcomes conditional on entry, fails to identify the distribution of fixed costs. First, entry decisions depend on multiple factors. We address this by jointly modeling fixed and marginal cost heterogeneity. Second, the average treatment effect does not trace out the full cost distribution, which is important for assessing surplus for inframarginal traders. This necessitates specifying a functional form for the marginal and fixed cost distributions. We assume that they are jointly lognormal:

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

This distributional assumption leaves us with 5 parameters to estimate: μ_{MC}, μ_{FC}, σ_{MC}, σ_{FC}, ρ_{MCFC}. We estimate the model via simulated method of moments with an optimal weighting matrix. As moments, we use entry probabilities by size of experimental subsidy and mean marginal cost intercepts times entry, by size of experimental subsidy. To estimate this second set of moments, we use our supply model (equation (21)) with estimated γ (Table 2) and an ω of 0 or 1, depending on whether the entrant has contacts, to back out marginal cost intercept estimates. For a candidate set of parameters, we draw potential entrants' marginal and fixed costs and then solve for the new market equilibrium under entry. We then compare variable profits, under entry, to fixed costs less the randomized entry subsidy to determine whether the potential entrant would choose to enter. We form the model moments from these simulated entry decisions. We use importance sampling to limit computational burden.

While our model contains distributional assumptions, the experimental moments prove highly valuable in identifying the cost distribution parameters. The randomization of the subsidy amount across ex ante identical potential entrants offers exogenous variation in net fixed cost that helps us trace out the shape of the fixed cost distribution. The marginal cost moments, which rely on our supply model estimates from

Table 7—Entry Model Moments and Estimates

Description		Estimate
Panel A. Model moments		
Weekly take-up rate - high offer		0.2982
Weekly take-up rate - medium offer		0.1789
Weekly take-up rate - low offer		0.0596
Entry × marginal cost intercept - high offer		3.76
Entry × marginal cost intercept - medium offer		2.59
Entry × marginal cost intercept - low offer		0.86
Marginal cost intercept - high offer		12.61
Marginal cost intercept - medium offer		14.49
Marginal cost intercept - low offer		14.43
Parameter	Estimate	Standard error
Panel B. Model estimates		
μ_{MC}	3.31	1.87
μ_{FC}	10.28	4.79
σ_{MC}	0.25	0.97
σ_{FC}	0.94	1.13
ρ_{MCFC}	0.14	0.73

Notes: Panel A presents the entry model moments and their sample values. The first six rows are the moments while the last three rows show the marginal cost intercepts conditional on entry implied by the moments. Note we use the unconditional moments in estimation but also present the conditional values here for clarity. Panel B presents the parameter estimates, with standard errors estimated from 1,000 bootstrap iterations. The parameters describe the log-normal multivariate distribution of marginal and fixed costs. Entry take-up rate is calculated weekly, such that each potential entrant contributes four observations.

Section V, are important because they help describe the relative importance of variation in marginal versus fixed costs in explaining entry decisions. Mean estimated marginal cost conditional on entry is similar across the three subsidy levels and, in fact, lowest for the high subsidy group (see panel A of Table 7). If marginal cost variation explained most entry choices, then we would expect that the low subsidy group would have the lowest marginal cost, conditional on choosing to enter. That we see little such variation implies that fixed costs drive most of the entry decisions.

Our entry model estimates, with standard errors, are in panel B of Table 7. The estimated parameters imply that, for the pool of potential entrants, mean marginal cost (intercepts) and fixed costs are about 28 Ksh and 46,000 Ksh, respectively, with a small positive correlation (0.14) between marginal and fixed costs. These means are unconditional and much higher than means conditional on entry. We estimate that entrants, in the absence of a subsidy, would have mean marginal cost intercepts of 19 Ksh and mean fixed costs of 17,000 Ksh. Fixed costs are high relative to variable profits and thus a meaningful consideration in welfare analysis. Our estimates are relatively imprecise, though as we will see in the next section, we still have enough precision to make meaningful welfare statements.

VII. Welfare and Counterfactuals

We now turn to estimating welfare and the division of surplus. With our estimates of demand and sellers' marginal costs, we only lack estimates of incumbent traders' fixed costs. In Section VI we estimated the fixed cost distribution for potential entrants. Applying this same fixed cost distribution to incumbents would ignore the likelihood that incumbents are not a random sample of potential entrants. We therefore make two refinements in estimating incumbents' fixed costs. First, all incumbents chose to enter, which means that variable profits must exceed fixed cost. Second, because marginal and fixed costs are correlated, we condition the fixed cost distribution on estimated marginal costs for incumbents. With these refinements, we estimate a trader's mean total profits in each market-week by integrating over possible fixed cost realizations.

We estimate that the mean (median) daily fixed cost for incumbents is roughly 4,200 (1,800) Ksh. For many incumbents, this cuts substantially into variable profits; for the median trader, fixed costs dissipate 71 percent of variable profits, leaving him with an estimated total daily profit of 1,200 Ksh. The median trader thus keeps 12 percent of revenues as profit.

We also calculate consumer surplus using our estimated demand system. We find that consumer surplus makes up only 18 percent of the total surplus remaining after accounting for deadweight loss, while traders reap 82 percent. While the standard errors on our entry model parameters are large, the 95 percent confidence interval for the consumers' fraction of surplus is $(15 \text{ percent}, 55 \text{ percent})$ and thus we can reject that consumers capture more than 55 percent of all surplus. In total, we estimate mean surplus per market-day as 64,000 Ksh (almost \$605).

We estimate that consumers capture only a small fraction of total surplus, and yet the typical trader keeps a fairly small fraction of revenue as profit. Underlying this, we find considerable heterogeneity in trader margins, as a long tail of traders sells very large quantities at high markups. Mean total daily profits are 13,000 Ksh (more than ten times the median) and 8 percent of traders see daily profits above 40,000 Ksh. The mean total profits, as a fraction of revenue, is 25 percent.[57] Further, because these higher markup traders tend to sell larger quantities, we find that the quantity-weighted mean total profits, as a fraction of revenue, is 45 percent. This heterogeneity implies that while the typical trader is not making large profits, the typical consumer is being served by a trader who is. Competition policy is therefore likely to be most successful if it is targeted, focusing on markets with these large, high-margin traders.

A. *Counterfactuals*

Even if competition policy could target the most profitable colluding traders and induce more competition, the gains to consumer surplus depend on trader heterogeneity and the shape of demand. We now explore these potential gains by conducting counterfactual exercises where we vary the level of competition (by varying ω). For each value of ω, we solve for a new quantity-setting equilibrium in each market-week; we start by keeping the set of traders fixed so that we can explore how variable

[57] Our estimates line up quite closely with existing estimates of profit margins collected from survey data with agricultural traders. For example, Fafchamps, Gabre-Madhin, and Minten (2005) find average margins of 11 percent and median margins of 8 percent in Benin, average margins of 27 percent and median margins of 11 percent in Madagascar, and average margins of 37 percent and median margins of 27 percent in Malawi. Other estimates of average margins range from 5–34 percent (see Dessalegn, Jayne, and Shaffer 1998 for estimates from Ethiopia; Minten and Kyle 1999 for the Democratic Republic of Congo; Gabre-Madhin 2001 for Ethiopia; and Fafchamps and Gabre-Madhin 2006 for Benin and Malawi).

surplus changes with competition. We remain agnostic about the actual policy that would induce more competition but rather focus on quantifying the potential gains.

Figure 9 shows the division of surplus as we move between Cournot competition and joint profit maximization. For joint profit maximization, which we estimate as describing current market conduct, consumer surplus is only 14 percent of total potential surplus, while traders' profits constitute 50 percent. The remaining 36 percent is deadweight loss from missing transactions that would have occurred had sellers priced at cost.[58] As we decrease ω to Cournot competition, consumer surplus rises while trader profits and deadweight loss fall. Under Cournot competition, the other well-defined form of conduct, we estimate that consumer surplus would increase to 42 percent, traders would capture 37 percent in profits, and deadweight loss would be 20 percent.[59] In terms of absolute magnitudes, this would represent an average increase of 36,000 Ksh ($333) in consumer surplus per market-day, and average decreases of 16,000 Ksh ($152) and 19,000 Ksh ($181) in trader profits and deadweight loss, respectively, per market-day.

The previous counterfactual ignores a potentially important margin: exit. If competition increases, then variable profits might fall by enough that some traders might exit (or fail to enter), which could potentially undo some of the gains from competition. We thus conduct a second counterfactual exercise where we impose Cournot competition but allow traders to exit if total profits are negative.[60] We find similar results to the counterfactual without exit. Consumer surplus is 6 percent lower and trader profits are 7 percent higher under exit, but these differences are small relative to the variation from the model of competition.

VIII. 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, limited trader record keeping and difficulties 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, we present some of the first experimental evidence on the topic. We implement trader cost-shock and demand subsidy experiments to estimate how traders compete within a model of supply and demand. We find evidence of a high degree of intermediary market power. Welfare analysis suggests that consumers enjoy only 18 percent of total surplus from these transactions, while intermediaries reap the rest. If traders priced at cost, total surplus would increase by 56 percent.

Given the high degree of market power observed, policymakers may be interested in pursuing policies that explicitly target enhanced competition among intermediaries. Our estimated counterfactuals indicate increased competition would yield large gains to consumers and improve market efficiency. However, antitrust

[58] These estimates differ slightly from the baseline welfare statements above because our counterfactual analysis includes a model of consumers arriving to the market. See online Appendix Section F for details.

[59] If traders priced at cost, consumer surplus would constitute 80 percent of the total. Traders would still earn profits because the marginal cost curve is upward-sloping.

[60] In some markets there could be multiple equilibria in terms of which traders exit. We solve for the equilibrium in which the traders whose profits would be most negative are the ones that exit.

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 may create additional problems, 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. In an additional experiment, we generate exogenous entry by offering traders subsidies to enter specific, randomly selected markets in which they have never worked before. We then estimate whether entry increases competition. We find that each additional trader entering the market reduces prices by only about 1 percent. Estimates suggest that traders with contacts are able to easily collude with incumbents and that those without contacts, though better able to encourage increased competition, are less likely to take up the entry offer in the first place. Thus, a broad policy of encouraging entry may find it challenging to generate the type of entry that increases competition.

Identifying mechanisms that increase competition is therefore an open challenge, given that collusive agreements seem flexible in incorporating the types of traders most likely to enter markets. 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 C. 2010. "Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger." *American Economic Journal: Applied Economics* 2 (3): 46–59.

Antràs, Pol, and Arnaud Costinot. 2011. "Intermediated Trade." *Quarterly Journal of Economics* 126 (3): 1319–74.

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 Paper 7179.

Ashenfelter, Orley, and Daniel Sullivan. 1987. "Nonparametric Tests of Market Structure: An Application to the Cigarette Industry." *Journal of Industrial Economics* 35 (4): 483–98.

Atkin, David, and Dave Donaldson. 2015. "Who's Getting Globalized? The Size and Implications of Intra-National Trade Costs." NBER Working Paper 21439.

Aw, Bee-Yan. 1993. "Price Discrimination and Markups in Export Markets." *Journal of Development Economics* 42 (2): 315–36.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2019a. "Common Ownership and Competition in the Ready-to-Eat Cereal Industry." Unpublished.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2019b. "The Common Ownership Hypothesis: Theory and Evidence." Unpublished.

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

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

Bergquist, Lauren Falcao. 2016. "Competitiveness among Intermediaries in Agricultural Markets." AEA RCT Registry. March 10. https://doi.org/10.1257/rct.682-3.0.

Bergquist, Lauren Falcao, and Michael Dinerstein. 2020. "Replication Data for: Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E119743V1.

Bergquist, Lauren Falcao, and Craig McIntosh. 2019. "Market Linkages for Smallholder Farmers in Uganda." Unpublished.

Berry, Steven T., and Philip A. Haile. 2014. "Identification in Differentiated Products Markets Using Market Level Data." *Econometrica* 82 (5): 1749–97.

Bresnahan, Timothy F. 1982. "The Oligopoly Solution Concept Is Identified." *Economics Letters* 10 (1–2): 87–92.

Bresnahan, Timothy F. 1987. "Competition and Collusion in the American Automobile Industry: The 1955 Price War." *Journal of Industrial Economics* 35 (4): 457–82.

Bulow, Jeremy I., and Paul Pfleiderer. 1983. "A Note on the Effect of Cost Changes on Prices." *Journal of Political Economy* 91 (1): 182–85.

Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel. 2019. "Selling Low and Buying High: An Arbitrage Puzzle in Kenyan Villages." *Quarterly Journal of Economics* 134 (2): 785–842.

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

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

Corts, Kenneth S. 1999. "Conduct Parameters and the Measurement of Market Power." *Journal of Econometrics* 88 (2): 227–50.

Crawford, Gregory S., Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu. 2018. "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *Econometrica* 86 (3): 891–954.

Dessalegn, Gebremeskel, Thomas S. Jayne, and James D. Shaffer. 1998. "Market Structure, Conduct, and Performance: Constraints of Performance of Ethiopian Grain Markets." Food Security Collaborative Working Paper 55597.

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

Fafchamps, Marcel, and Eleni Gabre-Madhin. 2001. "Agricultural Markets in Benin and Malawi: The Operation and Performance of Traders." World Bank Policy Research Working Paper 2734.

Fafchamps, Marcel, Eleni Gabre-Madhin, and Bart Minten. 2005. "Increasing Returns and Market Efficiency in Agricultural Trade." *Journal of Development Economics* 78 (2): 406–42.

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

Fan, Ying, and Christopher Sullivan. 2018. "Measuring Competition with a Flexible Supply Model." Unpublished.

Feenstra, Robert C. 1989. "Symmetric Pass-Through of Tariffs and Exchange Rates under Imperfect Competition: An Empirical Test." *Journal of International Economics* 27 (1–2): 25–45.

Gabre-Madhin, Eleni Z. 2001. *Market Institutions, Transaction Costs, and Social Capital in the Ethiopian Grain Market*. International Food Policy Research Institute Research Report 124.

Goldberg, Pinelopi Koujianou. 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." *Econometrica* 63 (4): 891–951.

Goldberg, Pinelopi Koujianou, and Michael M. Knetter. 1999. "Measuring the Intensity of Competition in Export Markets." *Journal of International Economics* 47 (1): 27–60.

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

Hjort, Jonas. 2014. "Ethnic Divisions and Production in Firms." *Quarterly Journal of Economics* 129 (4): 1899–1946.

Keniston, Daniel E. 2011. "Bargaining and Welfare: A Dynamic Structural Analysis." Unpublished.

Knetter, Michael M. 1989. "Price Discrimination by U.S. and German Exporters." *American Economic Review* 79 (1): 198–210.

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

Miller, Nathan H., and Matthew C. Weinberg. 2017. "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica* 85 (6): 1763–91.

Minten, Bart, and Steven Kyle. 1999. "The Effect of Distance and Road Quality on Food Collection, Marketing Margins, and Traders' Wages: Evidence from the Former Zaire." *Journal of Development Economics* 60 (2): 467–95.

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 2015-29.

Moser, Christine, Christopher Barrett, and Bart Minten. 2009. "Spatial Integration at Multiple Scales: Rice Markets in Madagascar." *Agricultural Economics* 40 (3): 281–94.

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica* 69 (2): 307–42.

Panzar, John C., and James N. Rosse. 1987. "Testing for 'Monopoly' Equilibrium." *Journal of Industrial Economics* 35 (4): 443–56.

Rashid, Shahidur, and Nicholas Minot. 2010. "Are Staple Food Markets in Africa Efficient? Spatial Price Analysis and Beyond." Paper presented at the COMESA Policy Seminar "Food Price Variability: Causes, Consequences, and Policy Options," Maputo, Mozambique.

Robinson, Amanda Lea. 2016. "Internal Borders: Ethnic-Based Market Segmentation in Malawi." *World Development* 87: 371–84.

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

Sullivan, Daniel. 1985. "Testing Hypotheses about Firm Behavior in the Cigarette Industry." *Journal of Political Economy* 93 (3): 586–98.

Sumner, Daniel A. 1981. "Measurement of Monopoly Behavior: An Application to the Cigarette Industry." *Journal of Political Economy* 89 (5): 1010–19.

Teravaninthorn, Supee, and Gael Raballand. 2009. *Transport Prices and Costs in Africa: A Review of the Main International Corridors*. Washington, DC: World Bank.

Villas-Boas, Sofia Berto. 2007. "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data." *Review of Economic Studies* 74 (2): 625–52.

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): 528–83.