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

Lauren Falcao Bergquist*

December 14, 2016

JOB MARKET PAPER
Find the most updated version **here**

Abstract

African agricultural markets are characterized by low revenues for smallholder farmers and high food prices for consumers. There has long been concern that this price wedge between farmers and consumers – and the resulting loss in producer and consumer welfare – are driven in part by imperfect competition among the intermediaries that connect them. In this paper, I implement three randomized control trials that are tightly linked to a model of market competition in order to estimate key parameters governing the competitive environment of Kenyan agricultural markets. First, I reduce the marginal costs of traders in randomly selected markets, and find that only 22% of this cost reduction is passed through to consumers. Second, to elicit the shape of consumer demand that these traders face, I randomize price discounts and measure the quantities that customers purchase at these prices. Taken together, these estimates reveal a high degree of collusion among intermediaries, with large implied losses to consumer welfare and overall market efficiency. Third, given that a natural policy response to limited competition is to encourage greater firm entry, I randomly incentivize the entry of new traders into markets. By capturing the resulting effect on local market prices, I identify the implied change in the competitive environment due to entry. I find that the estimates are consistent with a model in which entrants are able to easily join collusive agreements with incumbents. Taken together, the results suggest that agricultural traders in Kenya have considerable market power, and that marginal changes in market entry are unlikely to induce significant changes in competition. These findings have implications for the incidence of technological and infrastructure changes in African agriculture and for the policy responses aimed at improving the market environment.

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

*Department of Economics, UC Berkeley. Email: lfalcao@econ.berkeley.edu. I thank Edward Miguel, Marshall Burke, Benjamin Faber, Benjamin Handel, Kaushik Krishnan, Nicolas Li, Jeremy Magruder, Jamie McCasland, Craig McIntosh, Odyssia Ng, Katalin Springel, and Michael Walker for useful discussions. I also thank Meshack Odhiambo Okello, Ben Wekesa, Deborrah Muthoki Wambua, and Innovations for Poverty Action for excellent research assistance in the field. This research was supported by funding from the NSF Graduate Research Fellowship and the research initiative "Private Enterprise Development in Low-Income Countries" (PEDL), a program funded jointly by the 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. All errors are my own.

1 Introduction

The 1980s and 1990s saw a wave of liberalization sweep 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, agricultural markets remain poorly integrated, with prices varying widely across regions and seasons (Moser et al., 2009; Burke et al., 2016).

High transaction costs contribute to this limited market integration. Transport costs in Africa are the highest in the world (Teravaninthorn and Raballand, 2009); also prevalent are harder-to-measure costs associated with search (Aker, 2010), contractual risk (Fafchamps and Minten, 1999), and price uncertainty (Dillon and Dambro, 2016). However, much less is known about the degree of competition among intermediary traders in agricultural markets in developing countries. Whether traders are exerting market power matters for policymaking: if intermediaries are operating in a competitive environment in which price gaps are purely due to high transactions costs, then policies that reduce these transaction costs – road improvements, preferential terms for business expansion loans, and trade intelligence systems for broadcasting prices to traders, for example – would yield savings that traders will pass on to farmers (in the form of higher prices) and consumers (in the form of lower prices). On the other hand, if traders are exercising market power, gains from policies that reduce traders' operating costs may not be fully passed on to farmers and consumers; instead, the bulk of these benefits may be captured by intermediaries. To meaningfully improve farmer and consumer welfare in this environment, policies may need to explicitly target enhanced competition among intermediaries.

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

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

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

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

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

loss.

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

I use the model to solve for the predicted price changes resulting from entry under various counterfactuals for entrant behavior. Given estimated demand parameters, counterfactual simulations predict that the entry generated by the experiment should decrease prices by 8% if the entrant competes, 4% if conduct among traders remains unchanged, and 0% if the entrant colludes. This compares to the precisely estimated observed drop of 0.5% in the experiment, which is strong suggestive evidence of collusion. Structural estimates which jointly estimate demand and the change in the competitiveness parameter also find parameter estimates consistent with entrants colluding with incumbents. These results suggest that collusive agreements among intermediaries are flexible and can readily accommodate new entrants. Results from this paper therefore cast doubt on the power of entry by a small number of new traders to dramatically improve market competition in this setting.

This paper is one of the first to experimentally test the competitiveness of rural agricultural markets directly. Previous attempts to measure competition have mainly relied on observational methods. Observational studies have typically found high rates of pass-through across major markets (Rashid and Minot, 2010), though these high transmission rates may not extend beyond major urban markets (Moser et al., 2009; Fafchamps and Hill, 2008). Moreover, interpretation of this observational evidence is confounded by common shocks such as shared harvest times and reverse flows across seasons.[1] One exception to this primarily observational literature is a concurrent paper by Casaburi and Reed (2016), which studies the effect of an experimental subsidy per unit purchased

[1]Generally, it is difficult to cleanly trace price shocks from distinct points of origin to destination for ubiquitously produced and consumed staple agricultural commodities. Observational pass-through rates are much more informative for imported goods or manufacturing goods that have a distinct geographic point — and price — of origin.

to cocoa traders in Sierra Leone. They find small pass-through in terms of price, but larger pass-through in credit, suggesting the importance of interlinked relationships in their context (a feature not relevant in the Kenyan maize markets I study).[2] However, because their subsidy is offered only to a subset of traders in the market, Casaburi and Reed must ultimately rely on observational estimates of pass-through to measure the degree of competition, as their experimental estimates appear to be affected by within-market spillovers. Further, in the absence of evidence on the shape of farmer supply, they are forced to make strong linearity assumptions. Because the curvature of the market facing traders (farmer supply in their case, consumer demand in mine) is crucial to interpreting the pass-through rate and the implied degree of competition, I experimentally estimate this curvature.

Another set of papers attempts to directly measure traders' profits in order to draw inference about the size of rents and degree of competition.[3] These have generally found that average trader profits are high, though subject to large variability, leaving a question mark on whether these large returns represent rents or risk premia (Dillon and Dambro, 2016). Moreover, these direct measures are subject to severe measurement error in the face of difficult-to-quantify search, own labor, and fixed costs.(Fafchamps and Gabre-Madhin, 2006).[4]

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),[5] it is unclear if this suggests traders are already offering competitive prices given their transport costs or whether farmers are simply unable to utilize this information to improve their bargaining position.[6]

[2]The markets studied in this paper, in which traders sell to consumers, appear to more closely resemble spot markets. The vast majority of transactions – over 95% – are conducted in cash.

[3]In a more structural approach to firm-level data from traders in two markets in Ethiopia, Osborn (2005) finds in one market that traders pay a lower price when a larger number of farmers come to the market to sell, evidence of non-competitive pricing, and in the other no evidence of non-competitive pricing.

[4]Not to mention a general lack of record keeping among traders. In my sample, only 58% of traders keep any written records, and amongst this group, most records are fairly rudimentary. Directly asking about profits is exacerbated by the political sensitivity of the question, as traders explicitly fear being labeled exploitative.

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

[6]In a quasi-experimental variant of this literature, Casaburi et al. (2013) measure the impact of road expansion on market prices in Sierra Leone. While interpretation is complicated by the fact that the cost shock affects both farmers and traders, they conclude that the resulting price decreases can be best explained under a search cost framework, and are inconsistent with either Bertrand competition or Cournot oligopsony.

There is therefore a paucity of casually identified evidence on trader competitiveness (Dillon and Dambro, 2016), despite a growing interest in the role these intermediaries play in determining the allocation of gains from trade (Bardhan et al., 2013; Antras and Costinot, 2011).

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

This paper proceeds as follows: Section 2 describes the maize market industry in Kenya. Section 3 describes the theoretical model underpinning the experimental design, which is described in greater detail in Section 4. Section 5 presents results on pass-through, and Section 6 describes the demand estimation procedure. Section 7 presents the structural estimates of the level of competition among intermediaries and the welfare implications of these findings. Section 8 describes results from the entry experiment. Section 9 concludes and discusses policy implications.

2 Maize Markets in Kenya

2.1 Maize Output Market Chain

Figure 1 displays the maize output market chain in Western Kenya, as described by traders in interviews conducted by the author and in panel surveys conducted with over 300 regional traders in the area from 2013-2014.[7] Regional traders, the subjects of this study, are responsible for large-scale aggregation, storage, and transportation.[8] They report purchasing 50% of their maize from

[7]Note that the categories reported here on the trader sale side of the market deviate slightly from the categories utilized in the panel survey, as the survey was conducted during a pilot phase for this project, when the author was still in the process of identifying the most common and distant categories of actors to use in this framework. Therefore, the "local markets" category presented here is an amalgamation of a category entitled "retailers" and coded responses to an "other" category, which had a high number of "individual consumers in markets" listed. Further information on categorization utilized here is available from the author upon request.

[8]The definition of a regional trader employed in this paper is someone who both buys and sells in multiple markets.

small and medium farmers (selling less than 5 tons), 16% from large farmers, and 33% from other traders.[9] Traders tend to own a warehouse in a market center and either rent or own a lorry which they use to purchase maize, bring it back to their warehouse for sorting, drying, and re-packaging, and then carry onward to their destination of sale. In my sample, 64% of sales take place in open-air markets in rural communities. There, 66% of traders' customers are are individual households, while the rest are primarily village retailers. Traders also sell about 16% of their inventories to millers, who mill maize into flour for sale to supermarkets and other stores that serve urban consumers. They sell another 16% to other traders, who sell in other parts of Kenya or eastern Uganda. A very small portion of sales – about 2% – is sold to restaurants, schools, and other institutions. Finally, about 2% is sold to the Kenyan National Cereals and Produce Board (NCPB), 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.

2.2 Entry into Regional Trade

As part of a broad plan of structural adjustment in the 1980s and 1990s, Kenya pulled state-controlled marketing boards out of staple grain markets, lifted trade restrictions on export crops, and allowed prices to be determined by market forces, rather than by state mandate. Today, few legal barriers exist to entering into the maize trade.[10] However, engaging in large-scale, regional wholesale trade still requires significant working capital in order to pay for inventories, storage facilities,[11] and transport vehicles.[12] Further, traders must develop extensive networks of contacts

[9] About half of the purchases from farmers use a local assembler or broker. Brokers are often slightly wealthier members of rural communities (and are often farmers themselves) who identify other farmers in their villages who are ready to sell. They either purchase from fellow farmers and bulk to sell to the regional trader or, for a commission, they simply identify farmers who are willing to sell. Either way, they are small scale, often work only seasonally, and typically lack the working capital to do large-scale aggregation, long-run storage, or transport of any distance.

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

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

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

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.

Table 1 presents more details on these regional traders, the subjects of this study. The average trader has completed some secondary school (with 78% having completed primary school and 33% having completed secondary school) and is able to answer about half of the Ravens matrices Group B questions. Only 58% of traders keep written records, and those records typically only include prices and quantities of purchases and sales. Very rarely is cost or accounting data recorded. However, 62% do say they review their financial strength monthly. Most traders operate one-man businesses, with only 37% having any employees. Only 35% own their own truck; other traders rent means of transport.

2.3 Open Air Markets

This study takes place in the open air markets in which traders sell the majority of their produce.[13] These markets typically occur on a set day each week. Traders who sell in the market are a mix of those who have their warehouse in that particular market and those who arrive with a truck and sell out of the back for the day. Traders with lorries typically park next to each other in a particular area of the market that they use each week, and warehouses are typically in a row or cluster near each other. Importantly, prices set by other traders, while not posted in any public way, are presumably common knowledge given the close physical proximity of traders. Figure 2 presents the histogram of the number of traders per market, which varies from 1-10 with a median of 3. Traders commonly work in the same set of markets each week, with 95% of traders in my sample reporting working in that market most weeks and only 2% saying that this was their first time in the market (see Table 1). 77% have worked previously with *all* other traders in the market that day. As a result, 67% say they know other traders in their market on that day "very well", 27%

[13] Their customers in these markets are comprised of two-thirds individual household consumers and one-third rural retailers.

"somewhat well," and only 6% "not very well." When asked directly, only 38% of traders report "discussing a good price" with other traders and only 30% report engaging with other traders in an explicit agreement about the price at which they will sell; the vast majority claim they are rigorously competing on price. However, about 72% of traders work in a market in which at least one trader has reported to the existence of a price agreement that day.[14]

The median consumer buys maize only from his local market, though a few retailers purchase from a larger number. I therefore model consumers as being captive to their particular local market. The median customer buys maize for consumption every week; therefore, storage on the consumer side is rare.[15] The product itself is fairly homogenous.[16]

3 Theoretical Framework

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

3.1 Model Set-Up

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

[14]The discrepancy between self-reports cannot be explained by price agreements just among a small subset of traders in the market. When asked with *how many* traders they had this agreement, the average respondent who reports any agreement reports having this agreement with 77% of other traders (with 100% being the modal response).

[15]This data is drawn from a phone survey with 100 consumers randomly selected from the demand experiment sample. This survey was conducted in July and August 2016 immediately following data collection for the main experiment.

[16]See Appendix A for additional evidence on the extent of product differentiation.

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

Here, I employ a few simplifying assumptions. First, I assume that maize is a homogenous good and that traders are unable to price discriminate.[17] Taken together, these assumptions ensure that a single market price prevails and provide a theoretical link between market prices and individual traders' strategic interaction. Second, consistent with Fafchamps et al. (2005), I assume marginal costs c_{dt} are constant with respect to quantities. Finally, I assume symmetry across traders, specifically with respect to initial marginal cost. However, the feature crucial to the experimental design is that the *change* in costs is symmetric across traders; and the symmetric experimental manipulation of costs is explicitly designed to ensure this.

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

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

where Q_{dt} is the total quantity in the market, N_{dt} is the number of traders in the market, and θ is a "conduct parameter" $\equiv \frac{\partial Q}{\partial q}$ with the following interpretation:[18]

[17] There is little variation in quality and credit is rarely used. Further empirical evidence on product homogeneity is provided in Appendix A. The assumption of no price discrimination is based on the empirical context in which there is a high (0.9) intra-cluster correlation between the prices that a trader offers his various customers throughout the day. This is likely because negotiations between traders and consumers are conducted in public, thereby limiting traders' ability to engage in dramatic price discrimination (while there is no posted official price to ensure that prices are equivalent across customers, negotiations between traders and customers occur in front of the trader's truck or store, where other customers are typically lined up to purchase). However, traders may be able to engage in imperfect price discrimination using tools such as bulk quantity discounts, as documented in recent work by Attanasio and Pastorino (2015). In Appendix B, I explore this possibility further and find some limited evidence of imperfect price discrimination based on quantities. It should be noted, however, that any ability to price discriminate is prima facie evidence corroborating the existence of market power.

[18] The "conduct parameter" method has fallen out of favor in recent decades for many reasons (Corts, 1999), among which is that it only takes on a clear, well-defined interpretation at a few values. However, this formulation is convenient for the purposes of identifying predicted pass-through of cost shocks under a variety of standard models of market interaction and nests the primary models of concern, as described in this section. At values of θ other than these few well-defined values, one can treat it as a continuous heuristic for distance from perfect competition.

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

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

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

Sensibly, competitiveness in the market goes up with both the number of traders (holding conduct constant) and with more competitive conduct (holding the number of traders constant).[19] I summarize the competitiveness parameter under different models of competition:

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

Because σ synthesizes the components of competitiveness and yields a simple interpretation (the bigger σ, the higher the degree of competition), I will work with σ in the first portion of this paper, which measures the competitiveness of these markets. However, it is useful to keep the derivation of σ in mind when I turn to the effects of entry (i.e. increasing N) on conduct θ and overall competitiveness σ.

3.2 Pass-Through and Demand Curvature

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

[19] Recall here that a *smaller* θ represents more competitive conduct.

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

where $E_{dt} \equiv \left\{ \frac{Q_{dt}}{\frac{\partial P_{dt}}{\partial Q_{dt}}} \right\} \left\{ \frac{\partial \frac{\partial P_{dt}}{\partial Q_{dt}}}{\partial Q_{dt}} \right\}$ is the elasticity of the slope of inverse demand. Therefore, the level of pass-through ρ depends on both the competitive structure of markets σ and the curvature of demand E.[20]

Figure 3 provides a visual example of this relationship. In the left panel, a cartel determines how much of a Δ reduction in marginal cost c to pass-on to the price. With moderately curved demand, the cartel will chose to pass on only a fraction of the cost reduction. The right panel presents a market operating under Cournot competition but a more concave demand function. We see that a different combination of conduct and demand curvature could yield the same observable pass-through. Therefore, in order to infer conduct from pass-through, we must understand the curvature of demand.

3.3 Degree of Competition and Welfare Implications

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

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

[20] It is worth noting that, under this set-up, the prediction for pass-through under perfect collusion is observationally equivalent to under an alternative market structure in which traders sell perfectly differentiated products (i.e., when consumers' elasticity of substitution across products is zero). In the alternate structure, one could model traders as monopolists working in their own "markets" with an $N = 1$ and a $\theta = 1$ (and therefore $\sigma = 1$) as they are the only trader selling that particular type of good. If one assumes that E is the same across each traders' segment of consumers (a nontrivial assumption, but perhaps a reasonable first approximation as the elasticity of the slope of inverse demand is invariant to the size of the market), the pass-through rate would be the same as under perfect collusion, since σ and E would be identical under the two scenarios. Therefore, pass-through rates consistent with market power from perfect collusion are also consistent with market power from perfect product differentiation. Because it appears that maize is a homogeneous good in this context (see Appendix A), I interpret market power as arising from collusion; nonetheless, they are observationally equivalent.

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

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

where $\bar{\rho}$ is the quantity-weighted average pass-through rate.[21]

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

3.4 The Effect of Entry on Competition

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

4 Experimental Design

4.1 Sample Selection

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.[22]

[21] In Section 6, I describe (and provide empirical justification for) the functional form assumptions on demand that allow me to estimate a single ρ, rather than estimating a separate ρ at each quantity level.

[22] See Appendix C for additional details on the sample selection procedure.

4.2 Randomized Schedule

The two trader-level experiments (pass-though and entry) were each run for four weeks in a row. This time spans about a quarter of the full selling season in the region (March to July). This duration of treatment was selected to represent a long-run cost or entry shock. It was also selected to match the frequency at which prices regularly vary (see Figure 4) to minimize concerns about sticky prices.[23]

Because piloting revealed that market and week fixed effects were important (cutting standard errors almost in half), the experiment was designed to provide each market each treatment status (pass-through treatment, entry treatment, and control) in a random order. Figure 5 shows the six possible orders.[24,25] This allowed for the inclusion of market and week fixed effects in all analyses. Each four-week block was broken by a one-week break during which the demand experiment was run in a subset of markets.

4.3 Experiment 1: Pass-Through

In treatment market-days for the pass-through experiment, all traders in the market were offered a subsidy per kg sold. Enumerators arrived at the market at 7:30am (prior to the approximate market start time) and immediately made the offer to every trader present. Any traders who arrived later were also presented with the offer immediately upon arrival. Enumerators stayed in the market until 5pm (after the approximate market end time). Maize sold during the enumerators' presence in the market was eligible for the subsidy.[26,27] When introducing the subsidy, enumerators first

[23] Appendix G presents the effects of each treatment, broken down by week.

[24] One additional benefit of this design is that it allows me to test for any long-run effects of the treatments, given the random ordering of the treatment statuses. I find no evidence of long-run effects of the pass-through experiment, but I do find evidence of some small long-run effects of the entry experiment, which are presented in Appendix F. Accounting for these long-run effects does not alter the overall implications of the results presented Section 8.

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

[26] Only maize sold in cash was eligible for the subsidy due to concerns about the ability of enumerators to verify the authenticity of credit sales. However, over 95% of sales are conducted in cash, so this restriction was often irrelevant.

[27] The subsidy was capped at the first 75 90kg bags sold to limit budget exposure, but this cap was binding for only 1.5% of traders.

asked the trader to describe some of the major costs that he faced in his business.[28] 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 of the subsidy, traders were informed that this offer would be available for four weeks. An identical script was read in each subsequent week to remind returning traders of the availability of the subsidy and to make the offer to any new traders who were absent in the previous week. Therefore, all traders in the market received an identical reduction in their marginal costs, a crucial feature to map the experiment to the theory of pass-through.

The 60 markets in the sample were divided into two groups: 45 markets received a "low" subsidy level of 200kg/90kg bag when they were in the pass-through market treatment and 15 markets received a "high" subsidy level of 400kg/90kg bag.[29] Note that "low" and "high" are merely relative titles: both represent large and meaningful changes to traders' costs. The "low" subsidy rate represents 7.5% of the average price, while the "high" subsidy represents 15% of the average price. Payments were made via M-Pesa twice a day.[30]

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 monitoring of transactions and general data collection process was identical in treatment and control markets.

4.4 Experiment 2: Demand Experiment

In the demand experiment, customers were first allowed to approach traders and negotiated a price and quantity in a natural way before being approached by an enumerator to invite them to the demand experiment. If the customer gave consent, a random discount amount was drawn (using a randomization feature with SurveyCTO, the software utilized for electronic data collection) 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

[28] Trader in control market days were also asked these questions, to avoid confounding treatment with any priming effects.

[29] Sales of partial bags were eligible at the same prorated amount.

[30] Rapid payment was done in order to build among traders. Payments were made at 12pm and 5pm.

in light of this new price. The price discount was given to the customer in the form of an M-Pesa or a cash transfer[31] and the customer paid the trader the originally negotiated price.

Traders' consent was acquired at the beginning of each day and therefore the trader was aware that his customers would (potentially) receive a price reduction. While this may have changed the baseline price charged by the trader (for example, it is possible that the trader raised the average price to customer in order to collect some of the anticipated discount), the trader did not know at the time of price negation the actual amount of the discount that would be offered to the customer in question nor was the trader permitted to adjust the price following the announcement of the realized discount amount. Therefore, any variation in the price driven by the discount is random.

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

4.5 Experiment 3: Entry Experiment

In the entry experiment, traders who had never before worked in the treated market were offered subsidies to enter the market and attempt to sell there. For each market, three traders were given the offer. This was designed (1) to increase the probability that at least one trader took up the offer and (2) to measure traders' willingness to enter, as the amount of each offer was randomized. Offers were given for four weeks in a row, in order to generate somewhat long-run entry.

The pool of traders eligible to receive the entry experiment offers was drawn from the sample of traders interviewed in previous pilot work for this study (traders generally from markets in this 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 lorries were then excluded from the pool as pilot work showed that these traders almost categorically did not take up the offer. A phone survey was conducted of the remaining 187 traders to determine markets in which they had ever worked. For each of the 60 sample markets, I then identified the set of eligible traders who (1) had never before

[31] M-Pesa is the mobile money payment system common in Kenya. The customer was able to choose the method of payment.

worked in that market and (2) did not work in other study markets that occur on the same day of the week in order to avoid inducing exit in our sample. The median market had 37 eligible traders, the minimum had 28, and the maximum had 56.

From each of these sets, I then randomly selected the three traders who would receive the entry offers. Because I did not want to overwhelm a single trader with too many offers (potentially to the detriment of total take-up), I only offered each trader one offer per 4-week block.[32] Because this has cascading effects for the set of eligibles for each market, I randomize the order in which markets were assigned three traders from their remaining pool.[33]

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 (about $50 USD), a "medium" offer of 10,000Ksh (about $100 USD), and a "high" offer of 15,000Ksh (about $150 USD). The trader was eligible to receive this amount each time s/he visited the particular entry market in question on any of four offer days.[34] Payout was contingent on a few factors, all of which traders were made aware during the offer call. They were: (1) must come to the specified market on the specified date; (2) must arrive with a lorry and at least 15 bags; (3) must stay for at least one hour and show intention to attempt sales. Payments were made via M-Pesa immediately after these conditions had been met.

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

[32] In practice, this meant that once a trader was selected to receive an offer to enter one market, he became ineligible to be offered the entry incentive for any other market in the same 4-week block.

[33] In the first block, a few traders asked to be removed from the study (due to lack of interest in the subsidy and therefore unwillingness to answer surveys). When these traders were scheduled to receive an offer in a subsequent four-week block, they were then replaced and the offer was given to a new, unassigned trader from the same pool.

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

4.6 Data

Data was collected in an identical way in all markets and in all periods (pass-through treatment, entry treatment, and control). Depending on the activity level of each market, enumerators were assigned to survey 1–4 traders.[35] These surveys captured transaction-level price, quantity, payment method (cash or credit), and customer type (individual household consumer or retailer). I also collected data on the value of any non-traditional reductions in price; 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 service," free sacks, transport, or services given to customers bundled with their transactions.[36] Maize quality data was also collected for each trader.[37] Enumerators were trained to grade quality on an objective scale from 1 (lowest quality) to 4 (highest quality).[38]

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.[39] Finally, the first time a trader was met in the sample, additional information was captured

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

[36] These non-traditional reductions in price were not very common, but they do add 1–2 percentage points to my measure of pass-through, so there is some indication that traders can use these less-traditional methods of price reductions to pass-through some of the cost reduction. It is possible that this is a more discrete method of deviating from price agreements maintained with fellow traders. Alternatively, they may be used as tools of price discrimination between customers.

[37] There is no variation in quality offered by a single trader to his customers in the same market-day. In fact, traders commonly mix bags they have purchased of different quality prior to arrival at the market with the explicitly goal of offering a uniform quality level. I therefore collect only one measurement of quality for each trader in each market-day.

[38] The use of grain standards in Kenya are restricted to the most formal settings of large millers and the National Cereals and Produce Board. Regional traders typically do not know the official grade of their maize, and consumers do not use grades to describe or evaluate quality. Instead, traders and consumers assess quality of maize based on several readily observable characteristics: coloration, grain size, grain intactness, presence of foreign matter, and presence of weevil infestations. The following standards were developed with the help of several traders from the pilot: 4=Excellent [No pest,No foreign matter, No broken grain, No discoloration, sizable grain]; 3=Good [Barely infested, <5% foreign matter(e.g. maize cobs, dust, sand etc), <5% broken grain, <5% discolored]; 2=Fair [Infested, 5%-25% foreign matter, 5%-25% broken grain, 5%-25% discolored]; 1=Poor [Infested, >25% foreign matter, >25% broken grain, >25% discolored]. No formal tools were used to measure precise percentages; rather, enumerators were trained to take a handful of maize in their palm and count the kernels that matched each description (while this involves some imprecision, it is nearly identical to the process by which consumers judge quality — that is, by feel, sight, etc. — and therefore captures well the information available to consumers, which is the pertinent metric). Enumerator training included practice evaluating the quality level of real samples of maize.

[39] Note that, due to the sensitivity of these questions, these questions were asked at mid-day, after the enumerator

on the trader's fixed characteristics, including ethnicity, location of home market, highest level of education achieved, and a battery of business management and record keeping questions drawn from McKenzie and Woodruff (2015). A Raven's test was also administered.

My primary outcome of interest – price – is defined as the quantity-weighted average of transaction level prices that the trader sold that day.[40,41,42]:

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

where p_{idwt} is the price of transaction t for trader i in market d in week w and q_{idwt} the quantity.

This measure of price per trader P_{idw} forms the basis for the primary analyses of the pass-through and entry experiments.[43] All estimates are weighted by the inverse of the number of traders in the market so as to give equal weight to each market in the final analysis. All standard errors are clustered at the level of the market-block, the unit of randomization.

5 Pass-Through

This section presents the results on pass-through from the first experiment.

5.1 Estimation

To measure pass-through, I estimate:

had established good rapport with the respondent. However, there were several traders who left the market before that time. For these traders, enumerators attempted to ask these questions before the trader left but often failed given the short time between being alerted to their imminent departure and actual departure. As a result, there is higher "attrition" among this section of the survey.

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

[41]When non-traditional reductions in price were offered, the value of these additional services were captured and incorporated into the price per kg.

[42]Maize in these markets is traditionally sold either in 90kg bags or in 2.2kg "goros." Though there is some variation in the number of kgs contained in each of these units, I use these standard measure of 90kgs and 2.2kgs respectively, because piloting showed that variation was small and traders are often accused of lying about the number of kgs per unit, making self-reports of quantity unreliable.

[43]This is both in order to match up with the actual unit of observation in surveys – the trader – as well to line up with theory, which makes predictions about prices and quantities at the individual firm level; moreover, this is the unit of observation in surveys.

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

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

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

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

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

5.2 Results

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

I explore heterogeneity by the number of traders in the market.[44,45] Figure 6 presents these

[44]This is the main source of heterogeneity pre-specified in a design registry submitted prior to the beginning on the experiment. The design registry for this experiment can be found **here**, and was registered on March 10th 2016.

[45]The number of traders is defined as the average number of traders observed in the market over the course of the experiment. In order to remove any increases in the number of traders driven by the entry experiment, this figure

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

I further explore other sensible dimensions of heterogeneity by a few other measures in Figure 7.[46] First, I measure whether pass-through is different for markets on and off tarmac roads, which serve as a proxy for market geographic isolation. I find no evidence of heterogeneity by this measure. Next, I explore whether a higher intensity of explicit collusion predicts lower pass-through rates, measured by looking at the number of market-days within a market where traders have explicitly admitted to collusion.[47] The point estimates suggest that pass-through is sensibly smaller for markets above the median in this measure, but these differences are neither statistically significant nor large in magnitude. In summary, the lack of clear heterogeneity and relatively consistent point estimates suggests that pass-through is fairly constant across markets.

6 Demand Estimation

As described in Section 3, in order to draw inference about the level of competition from the observed pass-through, one must first understand the curvature of demand. I do this using results from the demand experiment.

6.1 Constant Pass-through Demand Class

I estimate a very general Bulow-Pfleiderer class of demand functions:

uses the average of the predicted number of traders each week, based on market and week fixed effects.

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

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

$$Q_{dt}(P_{dt}) = \begin{cases} \left(\frac{a-P_{dt}}{b}\right)^{\frac{1}{\delta}} & if \;\; (P_{dt} \leq a, b > 0 \;\; and \;\; \delta > 0) \;\; or \;\; (P_{dt} \geq a, b < 0 \;\; and \;\; \delta < 0) \\ 0 & if \;\; P_{dt} > a, b > 0 \;\; and \;\; \delta > 0 \\ \infty & if \;\; P_{dt} \leq a, b < 0 \;\; and \;\; \delta < 0 \end{cases} \tag{12}$$

where $a \geq 0$

I choose this particular class of demand functions for its flexibility, tractability, and empirical foundation. First, this demand structure is flexible because it nests many of the functional forms common to the development and trade literature, including linear demand, quadratic demand, and isoelastic demand. The demand functional form is tractable because it produces a constant elasticity of the slope of inverse demand with respect to quantity (E) (Bulow and Pfleiderer, 1983). Recall that in the model described in section 3, the pass-through rate was determined by the competitiveness parameter σ and the slope of inverse demand E; these demand functions predict a constant pass-through rate, independent of the level of demand and the size of the cost shock. To see this, note that the inverse demand function is:

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

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

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

Equation 14 nicely predicts a constant level of pass-through for a given σ. Because markets are randomized into receiving the low vs. high subsidy rate, one can assume these two sets of markets have – on average – identical levels of competitiveness (σ). Under this demand function, one should expect to see identical pass-through rates across the two levels of cost reduction amounts (because these amounts were randomized, the expected σ is the same under each). This is consistent with

what we see in Column 2 of Table 2, which suggests remarkably similar pass-through rates for the two levels of cost reduction, lending empirical support to this choice of demand function classes to consider.

6.2 Estimation and Results

I utilize the randomized reduction in the price paid by consumers from the demand experiment as an instrument for price. The analysis is run with 1206 observations.[48] I estimate the vector of parameters $\Theta = (a, b, \delta)'$ in Equation 12 using generalized methods of moments with a vector of sample moments given by $m(\Theta) = Z'\xi(\Theta)$. Here, Z is a matrix of instruments formed by the stacked row vectors $Z_i \equiv (1, d_i, d_i^2)$, with d_i defined as the value of the discount amount randomly offered to customer i (recall d is one of the ten possible discount values). The vector ξ is the stacked residuals from a logged transformation of Equation 12 such that $\xi_i = logQ_i - \frac{1}{\delta}log(a - P_i) + \alpha$, where $\alpha \equiv \frac{1}{\delta}log(b)$. Thus, the parameter estimates are given by the GMM objective function:

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

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

Because there are two sets of possible constraints on the parameters in order to see positive, finite demand, I estimate the model under each set of constraints separately. I find that the minimand is smaller under the first set of constraints and so continue under this set of constraints.[49] Moreover, note that the second set of constraints, in which $\delta < 0$, would suggest pass-through rates of *greater* than 100% under imperfect competition, which is clearly inconsistent with what is observed in practice (though it is important to emphasize that the demand estimation here is in no way

[48] This is the sample of the demand experiment customers, trimmed to drop the bottom and top 2% outliers in terms of price and quantity to avoid undue influence of outliers.

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

constrained by the results by the pass-through experiment).[50]

Estimates are initialized at 500 randomly selected starting values, to ensure the minimization procedure does not obtain parameters for a local minimum[51]. I then generate bootstrapped confidence intervals by estimating these parameters on 1,000 random draws (with replacement) of the data.

Results are presented in Table 3, which show the point estimate and 95% confidence interval. Figure 8 presents the fit of the data, plotting the quantities chosen at each price as predicted by the subsidy. I should note that the confidence interval on δ is wide. The confidence interval around δ contains estimates ranging from near elastic demand ($\delta = 0.01$) to linear demand ($\delta = 1$) to very curved inverse demand ($\delta = 5.89$). This is because δ, which represents the elasticity of the slope of inverse demand (plus one), is a higher order object which I am underpowered to measure with great precision, even with over 1,200 observations from the demand experiment. However, we will see in the next section that even this limited precision is sufficient for our purposes. From the point estimate on δ, I can predict the level of pass-through that one should expect under various models of competition; I will find the prediction of one model to line up very closely with what is observed empirically. Moreover, even at the bounds of my estimate on δ, I can still reject that what I see empirically is consistent with other common models of competition.[52]

7 Degree of Competition and Welfare Implications

First, I demonstrate that the observed pass-through is very close to what the model would predict if traders are behaving collusively with a demand curve given by the parameter point estimates.

[50] Pass-through switches from less than to greater than 100% at $E = -1$ (in our functional form, $\delta = 0$), because this is where the slope of marginal revenue equals the slope of demand, such that the change in marginal costs is equivalent to the change in price (see Figure 3 for graphical intuition on this point).

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

[52] As those models would have required even *more* curved inverse demand to be consistent with the observed pass-through rate.

Given the point estimate on δ of 2.8, I use Equation 14 to estimate the average pass-through rate one should expect to see in my sample, under various models of competition. If markets are perfectly competitive ($\sigma = \infty$), we should see 100% pass-through. If markets are Cournot competitive ($\sigma = N$), we should see pass-through rates that vary with the number of traders: $\rho = \frac{N}{N+2.8}$.[53] Given the distribution of number of traders in each market, the expected pass-through rate in my sample should be 55% if markets are Cournot. Finally, if markets are collusive ($\sigma = 1$), we should expect to see 26% pass-through.

Figure 9 displays the bootstrapped distribution of ρ.[54] I see that the mass of the distribution of ρ is concentrated near the predicted pass-through of 26% under collusion.[55] The dotted lines, which identify the 90% confidence interval, clearly reject a ρ consistent with that predicted under a model of Cournot competition or perfect competition.

This figure does not take into account the fact that δ is estimated imprecisely. To account for this, I generate a bootstrapped distribution of σ by separately using 1,000 bootstrapped values of ρ and 1,000 bootstrapped estimates of δ.[56] Figure 10 presents this distribution, overlaid with the predicted σ under each model of competition.[57] I plot in red the value of σ predicted by the point estimates on ρ and δ. The point estimate of $\hat{\sigma}$ is 0.81, which is quite close to – and statistically indistinguishable from – the model benchmark of $\sigma = 1$ under perfect collusion.[58] Moreover, while the collusive market benchmark of $\sigma = 1$ lies squarely in the middle of the 90% confidence interval, the levels of σ predicted by a Cournot model and perfectly competitive model lie outside these

[53] This would predict that pass-through would be increasing in the number of traders. Note that we already saw in Section 5 that pass-through did not vary with the number of traders in way that is consistent with this predicted pattern.

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

[55] Note the the point estimate of 22.4% is slightly below the prediction of 26%. However, they are remarkably close in magnitude and are statistically indistinguishable. This minor discrepancy observed may be noise in the estimate of ρ, which would affect the point estimate of ρ, or noise in the estimate of δ, which would affect the ρ predicted under collusion. I account for the full extent of this noise in Figure 10.

[56] Note that since I estimate each parameter separately off of distinct experiments, there is zero covariance between these estimates, and therefore the random pairing that I do of the two bootstrapped estimates is appropriate.

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

[58] The minor deviation between the estimated value of σ and the collusive model prediction is not statistically significant and may be due to noise in estimates of ρ and, in particular, δ, which is measured imprecisely.

bounds and I am therefore able to reject them.[59,60]

7.1 Discussion

The observed pass-through rate is therefore consistent with an underlying market structure in which traders exert a high degree of market power. I can rule out with 90% confidence a Cournot competitive market or perfectly competitive market. Moreover, estimates of the competitiveness parameter are largely consistent with perfectly collusive markets. That said, there are other forms of market power that could also be consistent with the observed pass-through rate, such as perfect price discrimination.[61] While the weight of the evidence presented in Appendix A suggests that maize is a fairly homogenous good, results presented here cannot definitively differentiate between these two forms of market power. Nonetheless, maize is almost certainly not perfectly differentiated and therefore collusion likely explains some of the market power of the traders. The following section describes the welfare implications of this lack of competition.

7.2 Welfare Implications

What does this imply for the division of surplus between consumers and intermediaries? I use Equations 7 and 8 to solve for the ratios for consumer surplus (CS), intermediary surplus (IS), and deadweight loss (DWL).[62] Table 4 shows the results. At $\sigma = 1$ (the closest model-consistent value to the estimated σ of 0.81) and a ρ of 0.26 (that which would be consistent with this σ), I estimate that only 17.8% of the total surplus generated by the maize market accrues to consumers, while intermediaries reap 67.6% of the surplus. Another 14.6% is lost to DWL. Even at the upper limit of my estimate of σ (and the corresponding ρ), consumers are at most receiving 25.5% of the total surplus. Therefore, I see that intermediaries accumulate much of the gains from these transactions.

I can also conduct welfare counterfactuals by plugging into Equations 7 and 8 the value of σ that

[59] I am able to reject a Cournot competitive model is that 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% under a Cournot model, we would have required a δ of about 12.

[60] The σ of ∞ predicted by a perfectly competitive environment lies all the way to the right outside the range of the figure and is clearly rejected.

[61] As described in Section 3, perfect product discrimination would result in a σ of 1 as well.

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

corresponds to counterfactual forms of market conduct and the ρ that would be realized at each of these values of σ.[63] Table 5 presents the results of this exercise. I find that if instead markets were Cournot competitive, consumers would reap 49% of the total surplus, and of course if markets were perfectly competitive, they would receive 100%.[64] Part of this gain in consumer surplus is a transfer from intermediaries to consumers, but this may be in keeping with the preferences of a policymaker who places greater weight on the welfare of poor rural consumers rather than on intermediaries. Further, the reduction in deadweight loss is an objective improvement. Figure 11 presents the same results in a more continuous form.

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

8 Generating Entry

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

8.1 The Cost of Entry

In this context, the cost of entry is high. Because the offer amount is randomized, I can use traders' willingness to accept the offer as a measure of willingness to enter new markets. Table 6 presents take-up at each subsidy level (take-up defined as ever accepting any of the four market day offers). Sensibly, I see that take-up increases in the size of the subsidy: take-up is 12% for the low offer,

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

[64] This, along with the other welfare results, relies on the assumption of constant marginal costs. If marginal cost were increasing, intermediates would still reap some positive percentage of the surplus in a competitive environment (however, as documented in Section 3, empirical evidence is consistent with constant marginal costs). Further, if traders were to price at average cost under perfect competition, intermediates would also earn a positive percentage of the surplus equal in absolute magnitude to their fixed costs.

28% for the medium offer, and 42% for the high offer. However, this is in stark contrast to the percentage of traders who report that it would be profitable to take-up the subsidy given their offer size: 77% at the low offer, 80% at the medium offer, and 89% at the high offer.[65] Standing in stark contrast with the naive measurement of opportunity cost, the low take-up rate when presented with a profitable opportunity is suggestive evidence that some other unobserved cost prevents traders from entering new markets.[66]

I also look at heterogeneity in take-up by a few key variables pre-specified in the design registry. To do so, I estimate the following regression specification on the pool of 180 potential entrants:

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

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

Figure 12 displays the results. As presented earlier, a larger subsidy size increases take-up. Longer distances to travel are sensibly correlated with lower take-up; when comparing distance?s effect on take-up with the offer amount, an additional 50km in distance is roughly equivalent to

[65] In a survey conducted during the offer phone call, traders reported the revenues and costs they would expect to incur if they did not take up the offer and instead followed their typical schedule and similarly, the revenues and costs they would expect to incur if they did take up the offer. When the profits expected under take-up plus the offer amount exceed the profits expected under no take-up, I code this as reporting that take-up would be profitable. Reported here are revenues and costs for the full week surrounding the offer in order to capture any effect take-up might have on other sales in the week if travel time or inventory effects drive inter-temporal trade-offs. This choice is conservative; "reported profitable" rates are slightly higher if I use daily profits.

[66] Of course, low take-up could also be due to trader mistrust of the offer. Two factors make it unlikely that this concern played a major role. First, 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 survey and therefore was well-known by many of these traders. In fact, we had run a pilot in the previous year with identical offers that were given and paid out. Therefore, IPA had a reputation of following through on any monetary offers it made. Second, as I will show in a moment, factors like distance to the market and trader business size seem to play a role in take-up, suggestive that low take-up is due at least in part to expected factors, rather than universal mistrust of the offer.

a drop of \$46 USD in the offer amount.[67] Having contacts in the entry market is correlated with higher take-up (albeit not quite significantly). The value of having any contacts is equivalent to an increase in the offer amount of \$36, which is fairly large in magnitude. Being a large trader (in terms of being above median profits) is also correlated with higher take-up. The effect of having above median profits is large, and is equivalent to offering an additional \$52 to enter. Both the effects on having contacts and being above median profit 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. Finally, ethnic similarity does not appear to have any effect on entry.

Despite the low take-up per trader, because I made offers to three different traders per market, this offer generates a strong instrument for entry. 53% of all markets had at least one day (out of four) with entry. 38% of all market-days had entry. And 26% of all market days had more than one entrant. In total, an average entry market had an additional 0.6 traders present, an increase of 13% over the mean market size and 20% over the median.[68] I turn now to the effect of this entry on prices.

8.2 The Effect of Entry on Price

To measure the reduced form effect of entry, I estimate:

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

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

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

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

is, pass-through treatment periods are omitted). Under this specification, the coefficient of interest is β, which yields the percent reduction in price observed in the entry offer market.

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

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

in which $\hat{N}_{dw}$ represents the number of traders in the market that day, for which I instrument with the EOM_{dt} dummy.

Table 8 presents these results. Despite a strong first-stage effect on the number of traders, the reduced form effects are small and not quite significant, with only 0.6% drop in prices. The IV estimate suggests that the entry of one trader reduces pries by 1% (p-value of 0.101).

Figure 12 presents heterogeneity in entry effects along different dimensions of market characteristics, which should be interpreted with caution.[69] Panel A presents differences in take-up rates by markets with above vs. below median number of traders, markets on vs. off tarmac roads, markets with above vs. below median reports of price discussions, and markets with above vs. below median reports of price agreements. No clear differences in take-up are seen across these groups, with the exception of markets with a greater number of traders, which do have statistically significantly higher take-up. Panel B presents IV effects on price broken down by the same categories. No definitive patterns of heterogeneity emerge based on the number of traders in the markets at baseline (see Table 7 for further breakdown of these effects) or whether the market is on tarmac road. However, it does appear that what small decreases I do observe in price are concentrated in markets in which fewer traders report discussing or agreeing on price.[70]

[69] In addition to the usual caveats regarding the ability to draw inference on causality from heterogeneity analyses, here I add another: recall that the pool of potential entrants differs by market, necessitated by the requirement that one has never worked in that market before. Therefore, it is likely that variation in market characteristics are correlated with variation in characteristics of the entrant, and it is therefore difficult to separate what differences in observed effects are due to variation in market characteristics versus entrant characteristics. That said, from the policymaker's perspective, separating these two may not be crucial if the two are correlated in practice.

[70] While the IV effects are precise zeros in markets with above median reports of price discussions or agreements, they are -2.8% and -2.5% in markets with below median reports of price discussions or agreements, respectively (both significant at 99%).

8.3 The Effect of Entry on Competition

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

To explore this, I develop a few benchmark cases for how one would expect θ – and therefore σ and ultimately prices – to change with this degree of entry. I consider three possible scenarios:[71]

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

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

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

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

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

Assuming $\mathrm{E}[c_T - c_C] = 0$ (which is true by construction of the RCT, at least for incumbents, and empirically true for entrants as well):[72]

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

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

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

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

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

The observed price change therefore reflects underlying changes in the competitiveness parameter σ (as well as any shifts along the demand curve as prices move). I can therefore look at this problem in two ways.

Taking the point estimates on δ and a seriously, I can first evaluate how much one would expect prices to move with entry for the various potential expected σ_T. Table 9 presents these simulations. The top panel presents the simulated effect of entry on one trader (the instrumental variable effect) for each market size (as determined by the baseline number of traders). The middle panel presents the reduced form effects to be expected given the first stage effect (ΔN) observed for each market size. This is my preferred benchmark for the expected reduced form effects because the variation in the first-stage effects enters non-linearly into Equation 21.[73]

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

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

in which ΔP_S^{EB} is the change in price expected in a market of size S if entrants act according

[73]However, there is likely some noise in the first-stage take-up estimates (especially for cells with a smaller number of markets). For example, the point estimate on markets with 10 traders seems to suggest an increase in the number of traders that is outside the known bound of three, the total number of entrants given the entry offer. There is only one market in this bucket, and therefore this discrepancy is likely due to noise not absorbed by the market and week fixed effects. The other estimates, however, appear to be in a reasonable range. Nonetheless, I also explore using the average first stage effect of 0.58 for all markets as a robustness check to address this potential noise. This specification eliminates concerns about noise entering non-linearly into Equation 21. However, it also removes the real variation in ΔN that should enter non-linearly. Simulation results for this alternative specification are presented in the bottom panel of Table 9. I find that the average predicted price effects under this specification are actually larger than those predicted using per-market-size variation in take-up. This is because take-up is greater in large markets and the IV effect of entry is smaller in large markets. Therefore, using this alternative assumption would produce even large predictions for the price decrease one should expect to see under models in which conduct is unchanged or the entrant colludes. Because I am able to reject the smaller levels of price change predictions from the main specification, I am also able to reject the predictions under this alternative specification.

to $EB \in \{$*conduct unchanged, entrant competes, entrant colludes*$\}$. Therefore, $\overline{\Delta P^{EB}}$ identifies the average reduced form price effect one should expect to observe under each model of entrant behavior. I estimate this figure to be a decrease of 4% if conduct is unchanged, 7% if the entrant competes, and 0% if the entrant colludes with incumbents.

The reduced form effect observed of 0.5% is clearly closest to the scenario in which the entrant colludes. Figure 14 presents a graphical version of this intuition. The bootstrapped distribution of the reduced form effect on log price is shown.[74] Overlaid is the reduced form effects on log price that one would have expected if the entrant competes, if conduct remains unchanged, or if the entrant colludes. I observe that the mass of the distribution of the effects is only slightly to the right of what one would expect if the entrant colludes, and the 90% confidence interval can rule out alternative scenarios in which the entrant competes or even conduct remaining unchanged. This analysis, however, has used the point estimates of the demand parameters. Do we still have precision when taking into account the variance in these parameters?

My analysis suggests yes. Take Equation 21 and solve for σ_T:

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

I then sample from the entire dataset 1,000 times. For each sample, I estimate ρ and, using bootstrapped values of δ, estimate σ_C. I then estimate $E[P_T], E[P_C]$, and $E[P_T - P_C]$ for each sample. Finally, for each sample, I calculate σ_T using Equation 23. The kernel density of the resulting σ_C and σ_T are displayed in Figure 15. A Kolmogorov-Smirnov test cannot reject that these two distributions are the same (D=0.0183, p-val = 0.996). I therefore conclude that entry has left σ unchanged. More specifically, I can test to which scenario the change in σ most closely corresponds. Figure 16 presents these results visually, demonstrating that the mass of the change in σ lies at zero, lining up closely with the predictions if entrants collude. The 90% confidence intervals rule out conduct remaining unchanged or the entrant competing.

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

8.4 Discussion

The observed limited change in price therefore reflects that the entry generated had a negligible effect on competition. In an environment with a high degree of market power at baseline, this is consistent with entrants being able to easily join existing collusive agreements with incumbents upon arrival.[75] Further corroborating the interpretation of collusion is the fact that what little price decrease I observe is concentrated in markets with lower self-reports of collusion (see Figure 12) and in the final week of entry, suggestive of dynamic effects consistent with a breakdown in collusion in the final period (see Appendix G).[76]

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

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

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

[76] Appendix G provides details on the evolution of treatment effects across the four weeks of treatment. Column 4 of Table G.1 demonstrates that while entry by one trader had no significant effect on price during weeks 1-3, entry by one trader resulted in a price decrease of 2% (significant at 10%) in the final week 4. This is suggestive of some breakdown in collusion in the final stage of game, as predicted by theories of repeated interaction and dynamic collusion.

9 Conclusion

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

Taken together, these results suggests that policies commonly proposed to reduce the cost of agricultural trade – such as paving rural roads, implementing market price intelligence systems, and instituting uniform quality grading – would do little to achieve their stated aims of improving consumer and farmer welfare unless they also enhance the level of competition among traders.[77] Low pass-through rates indicate that traders retain the vast majority of reductions to their costs, rather than passing them on. Given the high degree of market power observed, policymakers may be interested in pursuing policies that explicitly target enhanced competition among intermediaries, which simulations indicate would yield large gains to consumers and improve market efficiency.

However, antitrust regulation of traders would probably be difficult to implement in an environment of low state capacity, and direct state intervention into the market to supplant the private sector would likely create more problems than it would solve, as seen during the largely unsuccessful experience with state-run markets following independence. Policies that encourage greater market

[77] Some policies to reduce traders' costs may indirectly enhance competition. For example, road construction could encourage entry of new traders into markets. However, I find here that entry does little to increase competition.

entry may be more a feasible response in this setting. However, I find that entry yields little benefits to consumers, at least at levels seen in this experiment. While it is possible that massive entry – for example, doubling the number of traders in a market – could do more to increase competition (the effects of such a treatment are outside the scope of this paper), evidence presented here does suggest that such a policy is at best likely to be expensive, given that willingness to enter new markets appears low among traders.

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

Fafchamps, Marcel and Bart Minten, "Relationships and Traders in Madagascar," *Journal of Development Studies*, 1999, *35* (6), 1–35.

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

_ **and Eleni Gabre-Madhin**, "Agricultural Markets in Benin and Malawi," *African Journal of Agricultural and Resource Economics*, 2006, *1* (1).

_ **and Ruth Hill**, "Price Transmission and Trader Entry in Domestic Commodity Markets," *Economic Development and Cultural Change*, 2008, *56*, 729–766.

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

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

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

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

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

Miguel, Edward and Michael Kremer, "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica*, 2004, *72* (1), 159–217.

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

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

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

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

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

Weyl, E. Glen and Michal Fabinger, "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, 2013, *121* (3), 528–583.

Tables and Figures

Figure 1: **Maize value chain in study area**. This figure displays the maize output market chain in Western Kenya, as described by traders in interviews by the author and in panel surveys conducted with over 300 regional traders in the area from 2013-2014. Note that the categories reported here on the sale side deviate slightly from the categories utilized in the panel survey, as the survey was conducted during a pilot phase for this project.

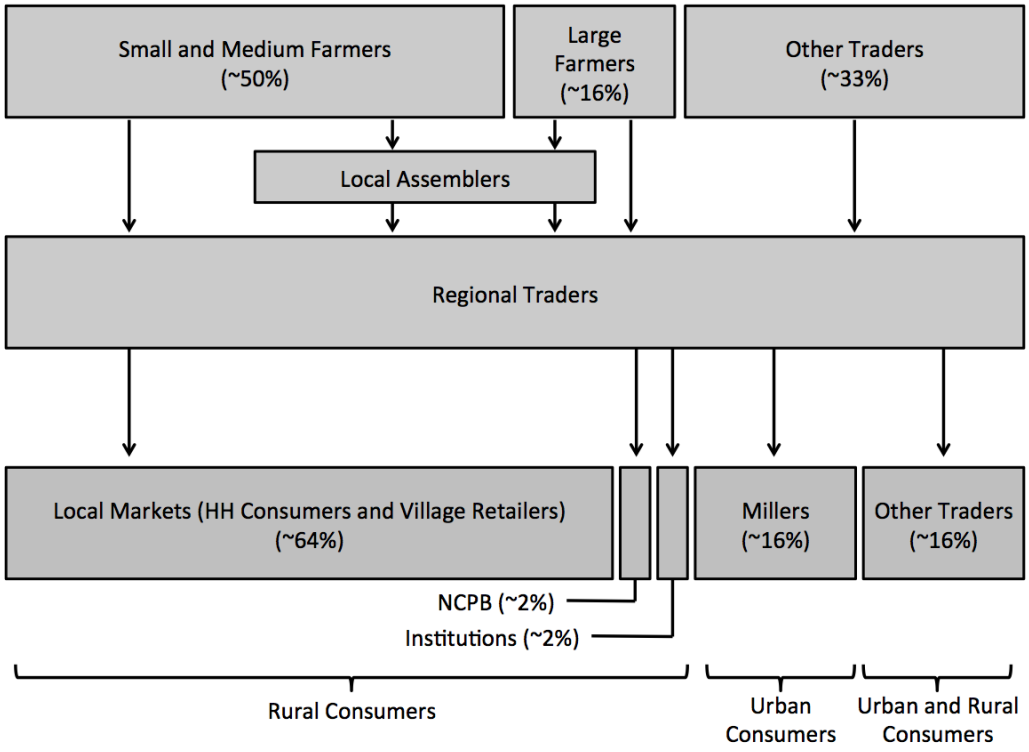

Figure 2: **Histogram of number of traders**. The number of traders is calculated as the average number of traders present in the market during 12 weeks of the study period, as predicted by week and market fixed effects (that is, any increase in number of traders due to the entry experiment is omitted).

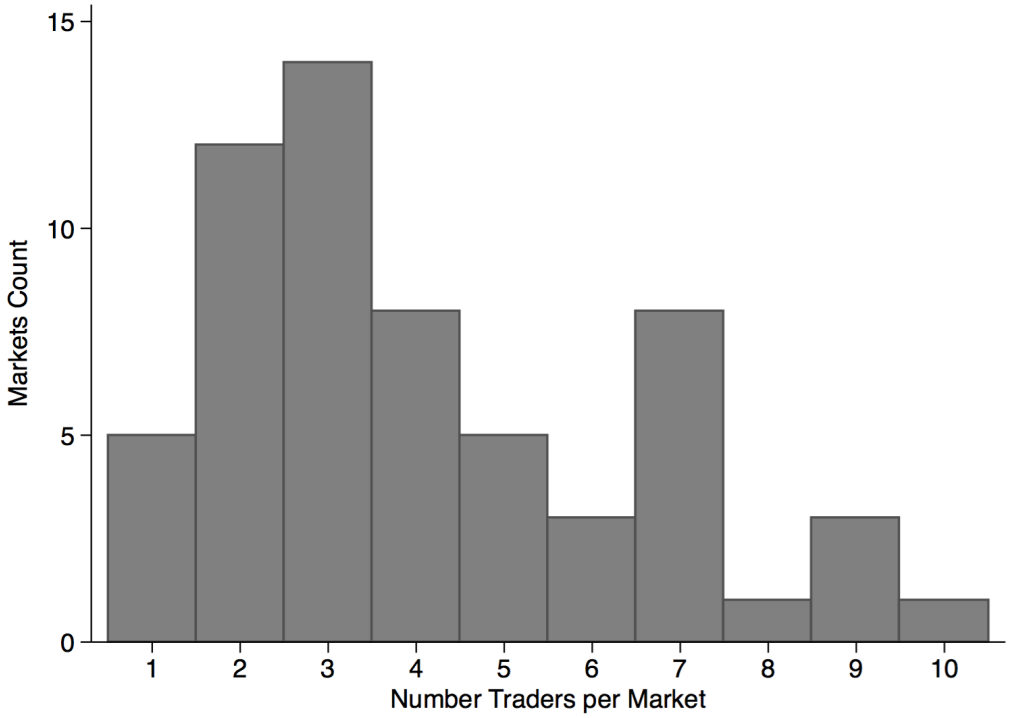

Figure 3: **Pass-through given conduct and demand curvature**. **Left panel**: Here I show a collusive environment, in which traders act as a single profit-maximizing firm. Quantities are set where the marginal revenue (MR) curve meets the marginal cost (MC) curve. Prices are then set where this optimal quantity intersects the demand (D) curve. A shift in the marginal cost curve downwards results in a pass-through rate of $\frac{\Delta P}{\Delta MC}$. With a concave demand function, as shown here, pass-through rates will be low. **Right panel**: Here I show a Cournot competitive environment, in which traders compete on quantities. The figure shows an individual firm's pricing and quantity decision. The firm takes the amount produced by other firms in the market (q') as given (producing a residual demand (RD) curve) and from this determines its marginal revenue curve. It then sets its own quantities where its marginal cost curve meets its marginal revenue curve. Prices are then set by where total quantities hit the demand curve. This figure illuminates how the same observable pass-through rate can be consistent with an underlying model in which markets are collusive and inverse demand is only somewhat curved, or an underlying model in which markets are Cournot competitive and inverse demand is strongly curved. I therefore must estimate the shape of demand to infer from pass-through the underlying market structure.

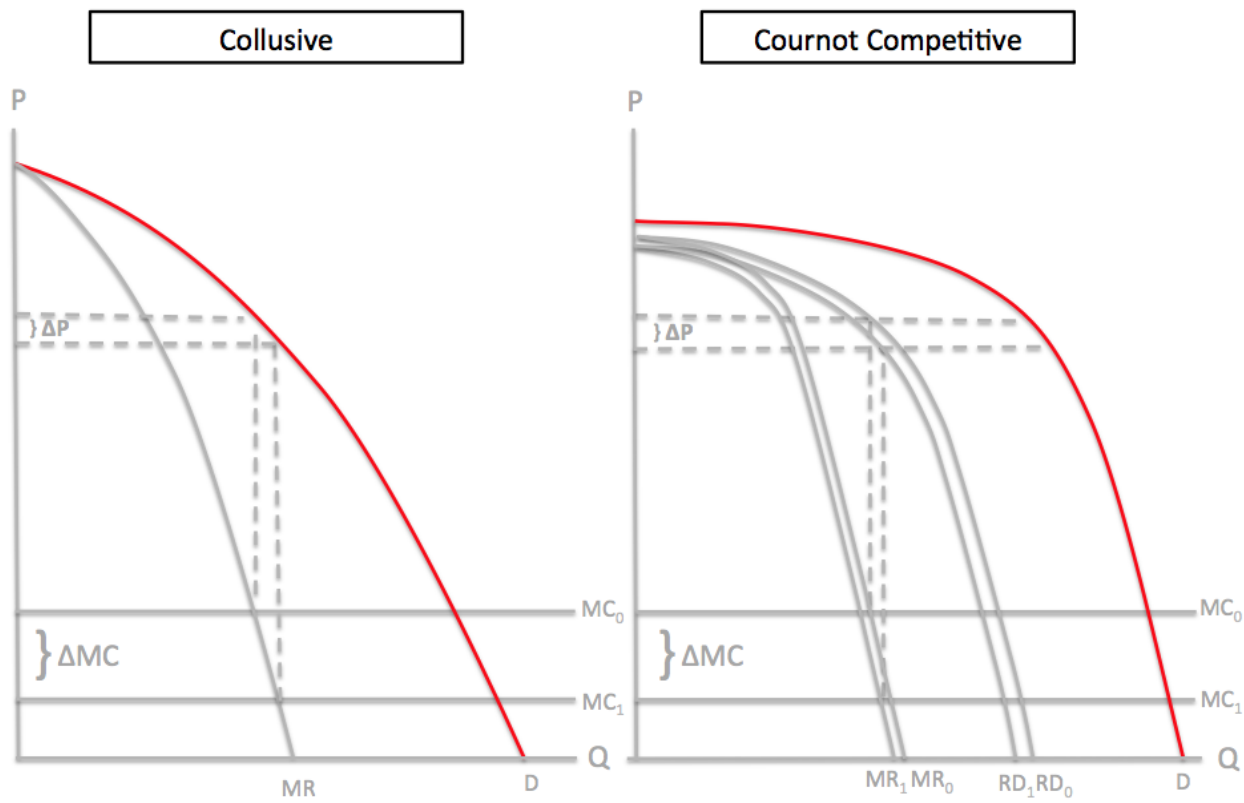

Figure 4: **Maize prices in study markets**. Grey lines show the price for each market over the 12 weeks of the study period. The black line shows the average price across markets in each week. The black bar on the vertical axis show the average size of the cost reduction subsidy for comparison. It appears that prices fluctuate by as much as the size of the subsidy each month (if not more for some of the more erratic markets). I view this as evidence against a story of sticky prices preventing greater pass-through.

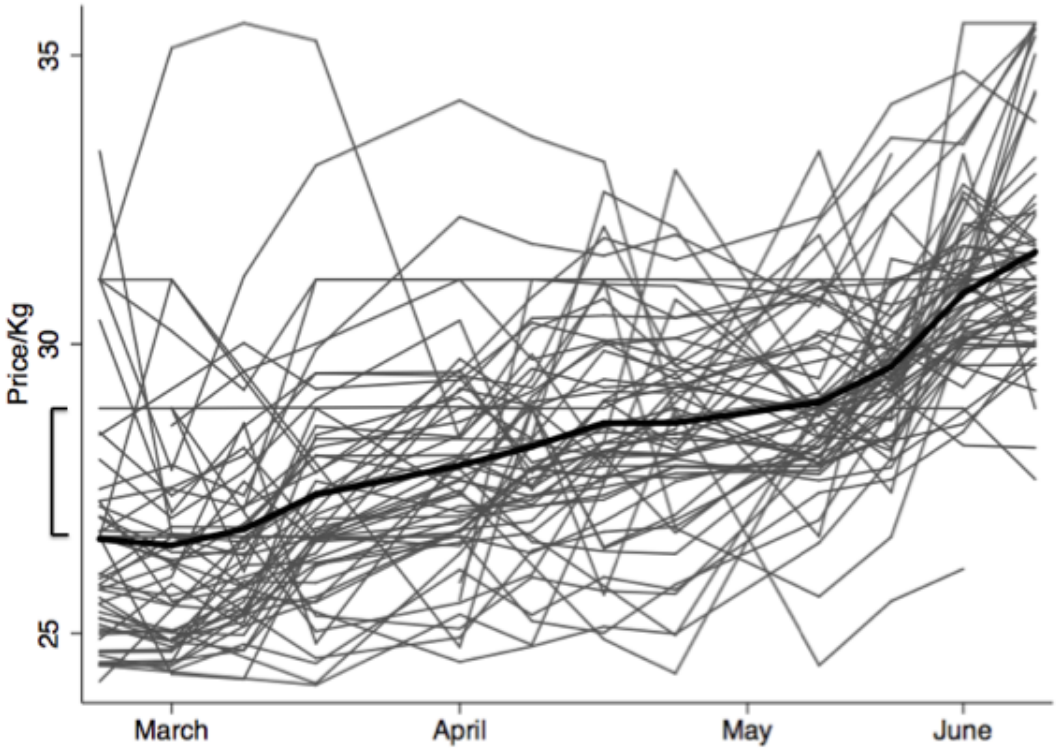

Figure 5: **Experimental schedule**. The 60 markets in my sample are randomly assigned one of six possible schedules, in order to yield randomized ordering of treatment statuses. There are therefore 10 markets in each schedule. This allows the inclusion of market and week fixed effects in every analysis. There is therefore a total of 720 market days in my sample, clustered into 180 market x four-week block cluster (standard errors in all specifications are clustered at this market x four-week block level). The demand experiment is run in a quarter of each markets during each week break in between each treatment status. Each market therefore receives the demand experiment once.

	Schedule 1	Schedule 2	Schedule 3	Schedule 4	Schedule 5	Schedule 6
Week 1	Demand Experiment in 1/4 of markets					
Week 2 Week 3 Week 4 Week 5	Pass Through	Control	Entry	Pass Through	Control	Entry
Week 6	Demand Experiment in 1/4 of markets					
Week 7 Week 8 Week 9 Week 10	Entry	Pass Through	Control	Control	Entry	Pass Through
Week 11	Demand Experiment in 1/4 of markets					
Week 12 Week 13 Week 14 Week 15	Control	Entry	Pass Through	Entry	Pass Through	Control
Week 16	Demand Experiment in 1/4 of markets					

Figure 6: **Heterogeneity in pass-through by number of traders**. No clear pattern of heterogeneity is seen in pass-through by number of traders. Most estimates are fairly closely clustered around the average pass-through rate of 22% (highlighted with the dotted line). The exception is the one market with 10 traders, but this appears to be a outlier driven by a small sample size. The bottom two estimates show pooled results, grouped into above and below median number of traders. Again, no heterogeneity is seen.

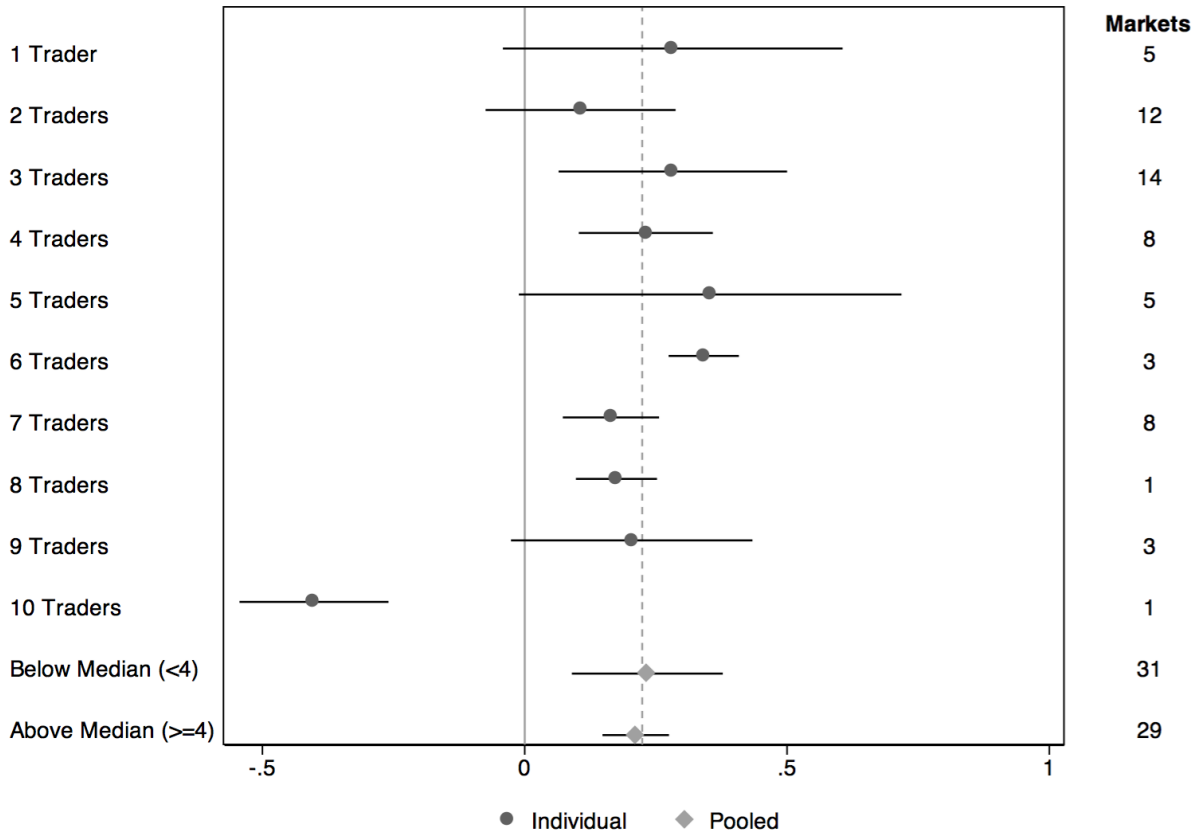

Figure 7: **Heterogeneity in pass-through by various factors**. No clear pattern of heterogeneity is seen in markets that are on tarmac roads (less geographically isolated) vs. off (more geographically isolated). The "discuss (agree) price" variable is constructed by counting the number of days in which at least one trader in that market stated that they had discsused (agreed on) a good price and dividing markets into above and below the median of this measure. Sensibly, the point estimates on pass-through are higher in markets in which discussing or agreeing to prices are more commonly self-reported, but this difference is not significant. It is unclear how much of this is due to measurement error due to self-reporting bias.

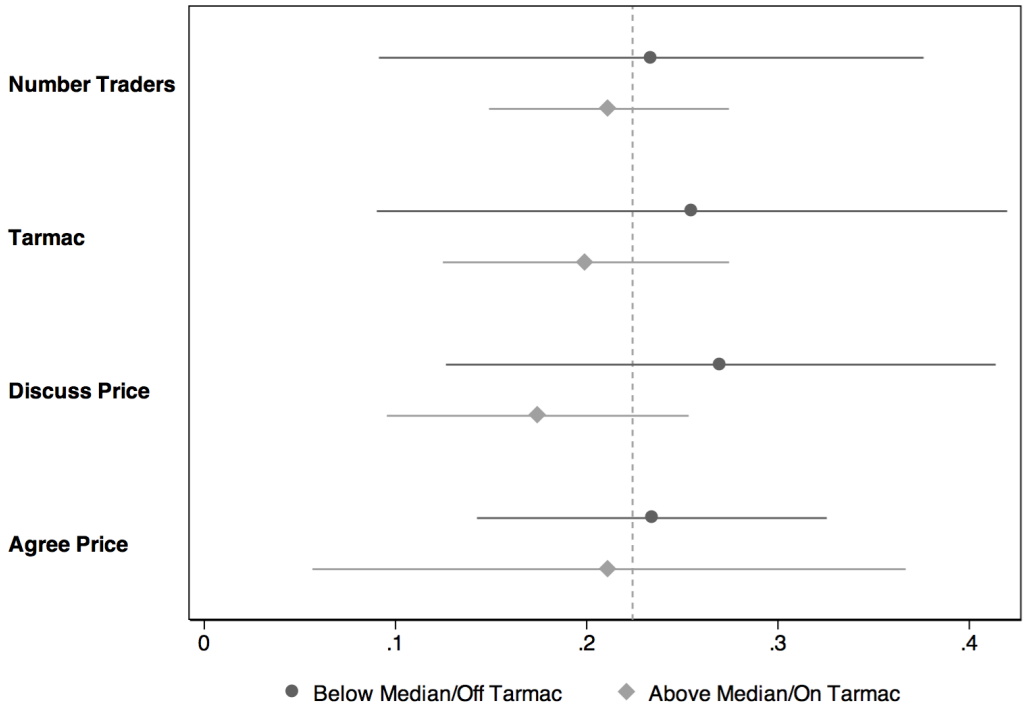

Figure 8: **Demand fit**. The x-axis plots the price as predicted by the (ten levels of) discounts offered in the demand experiment, while the y-axis plots the quantity purchased at that predicted price. The black line plots estimate Bulow-Pfleiderer demand functional form.

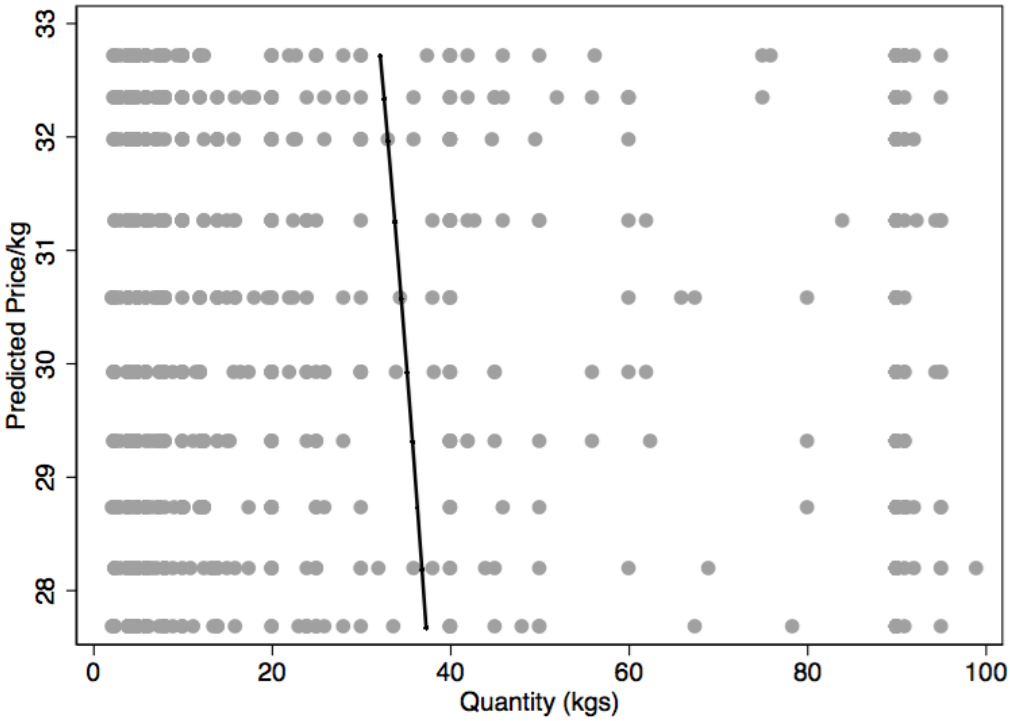

Figure 9: **Predicted pass-through under various competitive environments**. Given my estimate for the curvature of demand, I predict that one would have observed 100% pass-through in a perfectly competitive market, 55% pass-through in a Cournot competitive market, and 26% pass-through in a collusive market environment. Here I plot the distribution of pass-through, calculated by bootstrapping my primary pass-through regression 1,000 times. The bootstrapped 90% confidence intervals are shown in small dotted lines. The mass of the pass-through estimates line around the value predicted in a collusive environment, and I can rule out pass-through rates consistent with Cournot competitive or perfectly competitive markets with 90% confidence.

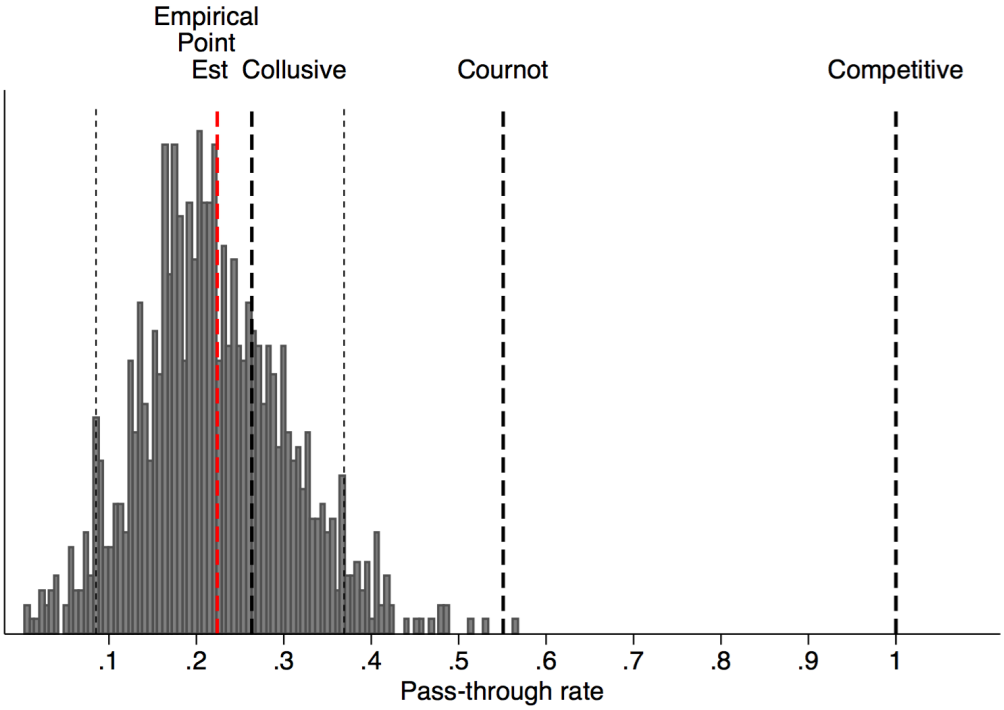

Figure 10: **Competitiveness parameter (sigma) estimates**. I combine my bootstrapped estimates of pass-through with my bootstrapped estimates of demand curvature in Equation 14 to generate 1,000 bootstrapped estimates of the competitiveness parameter σ. Recall that $\sigma = 1$ if competitive, N (the number of traders) if Cournot competitive, and ∞ if perfectly competitive. I find that the point estimate for σ is remarkably close to 1, the value predicted under collusion. Moreover, the 90% confidence interval can rule out a Cournot competitive environment (the value of σ plotted here under Cournot competition is the σ expected given the distribution of market sizes in my sample. Finally, I can clearly rule out perfect competition at $\sigma = \infty$, not shown here for obvious reasons.

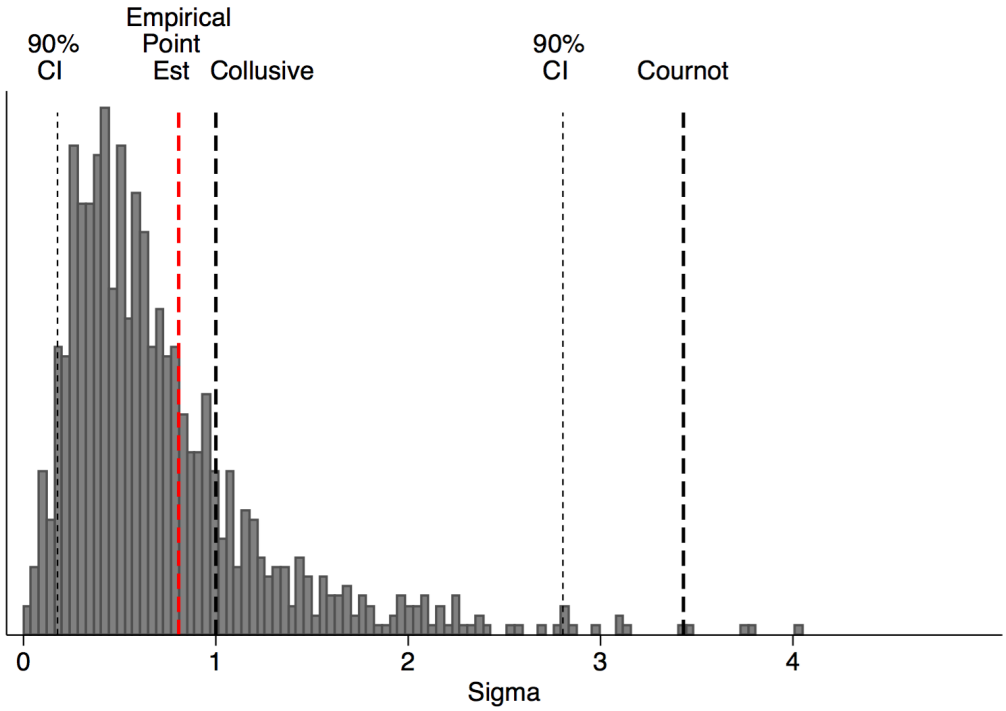

Figure 11: **Welfare counterfactuals**. Counterfactual division of welfare is shown for the average market of four traders. The current division of surplus is shown at the far left vertical dotted line, suggesting that intermediary surplus (IS) is 67.6% of total surplus, while consumer surplus (CS) is only 17.6% and deadweight loss (DWL) is another 14.6%. I see that moving to a Cournot competitive environment would yield clear improvements in welfare for consumers, as the division of surplus would be closer to 50/50 between consumers and intermediaries. In a perfectly competitive environment, consumers would reap all of the surplus.

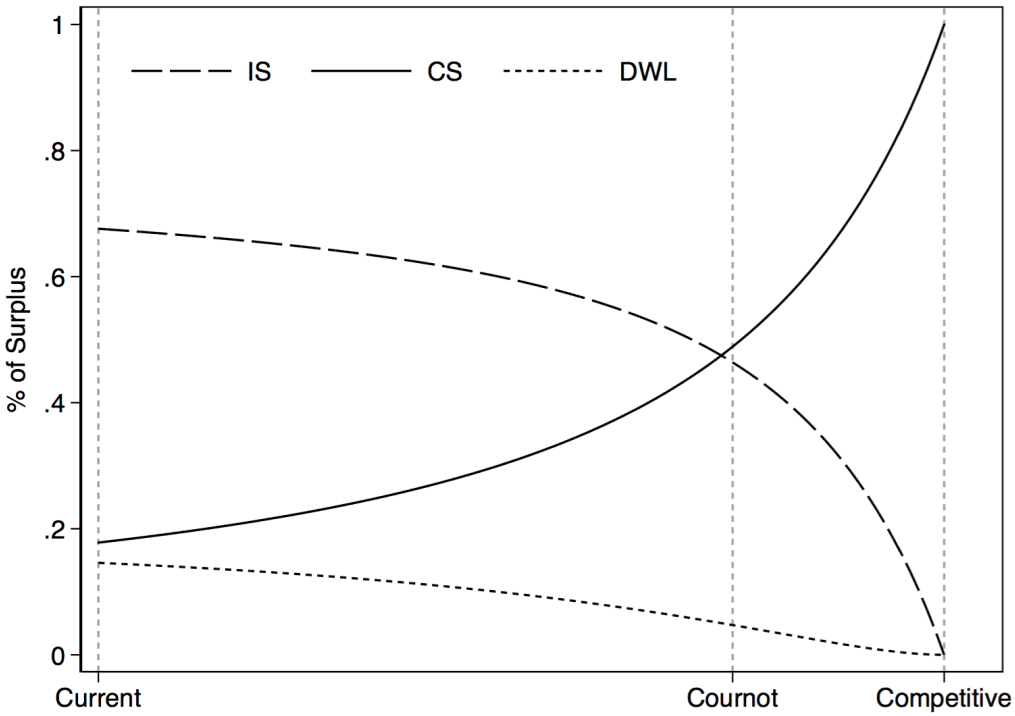

Figure 12: **Heterogeneity in willingness-to-enter**. Take-up of the entry offer is regressed on various measures of heterogeneity (alternately without and with market fixed effects; the latter compares only traders offered to attend the same market). The coefficient and 95% confidence interval is plotted. Take-up is sensibly higher for those with large offers. Distance is sensibly correlated with lower take-up (though this loses significance once market fixed-effects are included, likely due to loss of variation in distance).Having contacts and being a larger trading business (in terms of profits) is associated with higher take-up, suggestive that networks and working capital might be facilitate entry (and conversely, that a lack of these assets might serve as barriers to entry for others). Finally, perhaps surprisingly, traders are no more likely to accept an offer to enter markets in which a larger percentage of incumbents are of the same ethnicity.

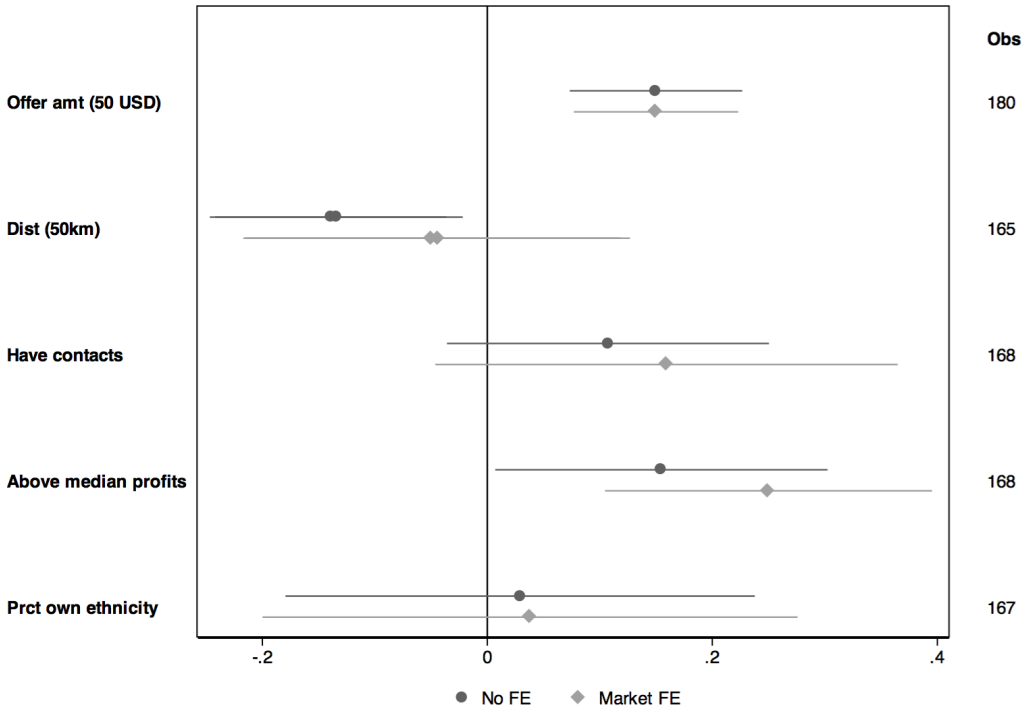

Figure 13: **Heterogeneity in take-up and IV impact of entry by market characteristics**. Number of traders divided into below and above median number (as counted in control periods). Tarmac is a dummy for whether the market is on a tarmac road or not. The "discuss (agree) price" variable is constructed by counting the number of days in which at least one trader in that market stated that they had discsused (agreed on) a good price and dividing markets into above and below the median of this measure. The unit of observation in Panel A is the market-day and the sample is restricted to entry treatment market days. Panel A presents the results from a t-test of a dummy for whether any entry occurred on that market-day by the relevant dummy. The mean and 95% confidence intervals for each subgroup are shown. Panel B uses the full trader sample and presents the point estimate and standard errors on an IV specification identical to that presented in Equation 18, but with the sample restricted to the subgroup in question.

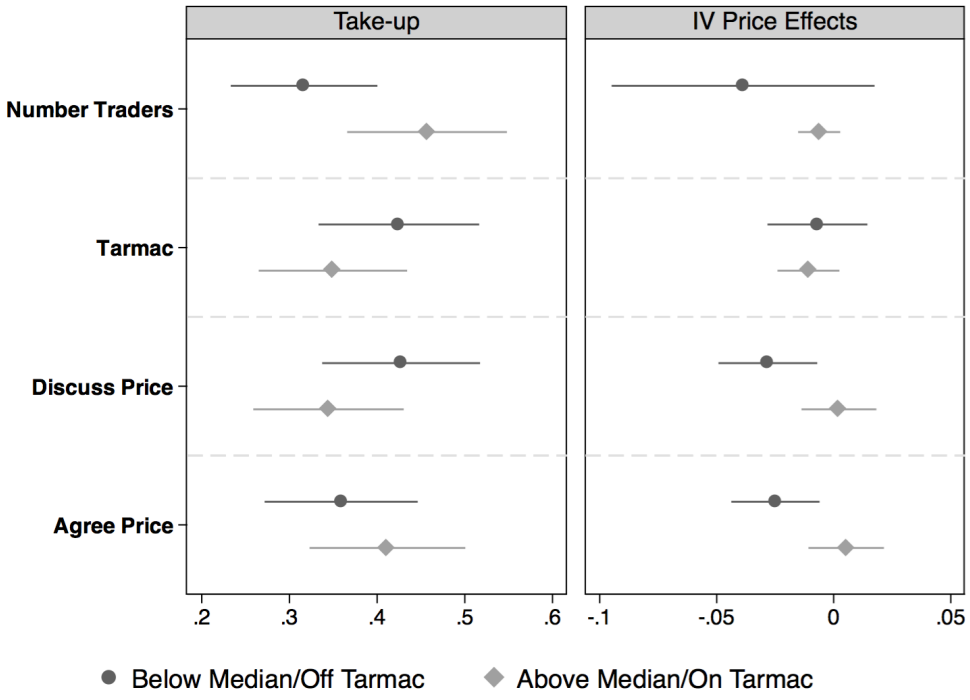

Figure 14: **Predicted price change for various formed of behavior on the part of the entrant**. Taking the point estimates on the demand estimation seriously, I plot the expected reduction in price one would have expected to see under different models of entrant behavior: entrant competing with the incumbent, overall market conduct unchanged upon entry, and entrant colluding with incumbents. I identify in vertical long-dash lines the price change one would have expected to see on average, given the distribution of baseline number of traders in my sample of markets, and given the first-stage effect of increasing the number of traders by 0.582. I overlay the distribution of my estimate of the price effect, estimated off of 1,000 bootstrapped samples of my data. I see that the mass of the price reduction seen lies fairly close to the 0% price effect predicted if the entrant colludes with incumbents, and I can rule out with 90% confidence that conduct was unchanged or that the entrant competed (90% confidence interval identified in short dashed vertical lines).

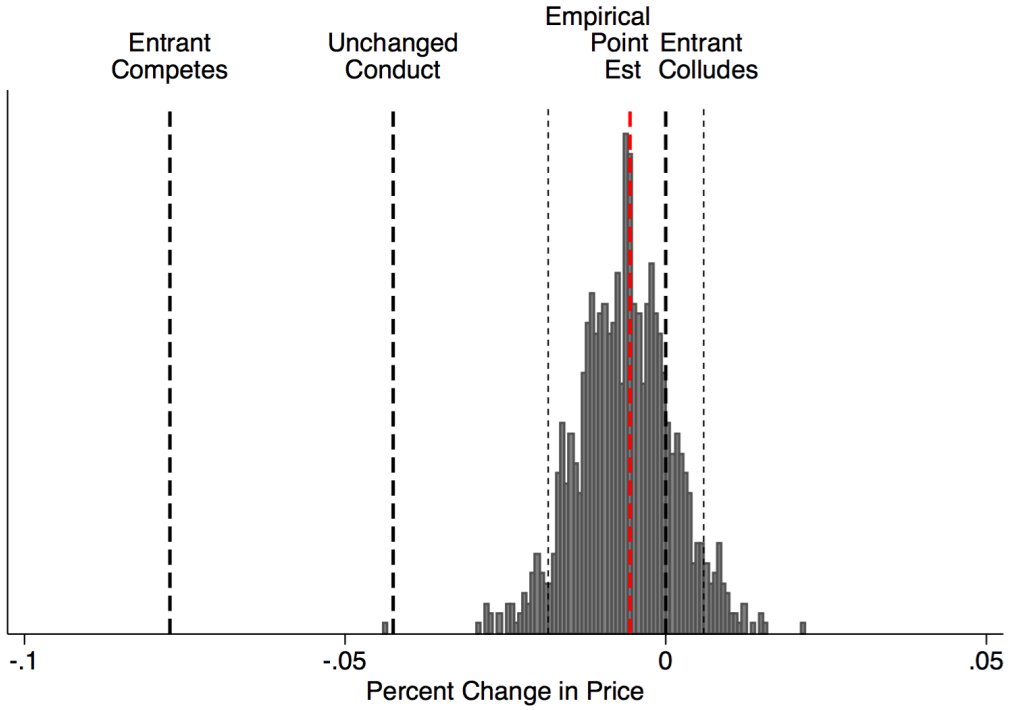

Figure 15: **Kernel density of sigma control and sigma after entry**. I take 1,000 bootstrapped samples of my entire data and, for each, calculate the baseline competitiveness parameter σ_C (as determined by the estimated pass-through rate int that sample and the estimate of demand curvature δ from the corresponding bootstrapped demand parameter estimates). I then estimate, for each sample, an estimate of σ_T under entry, using each sample's estimate of σ_C, the price effect of entry, and δ and a from the corresponding bootstrapped demand parameter estimates. The kernel densities of σ_C and σ_T are plotted here. Clearly, sigma has barely shifted right. Unsurprisingly, a Kolmogorov-Smirnov test cannot reject that these two distributions are the same (D=0.0183, p-val=0.996).

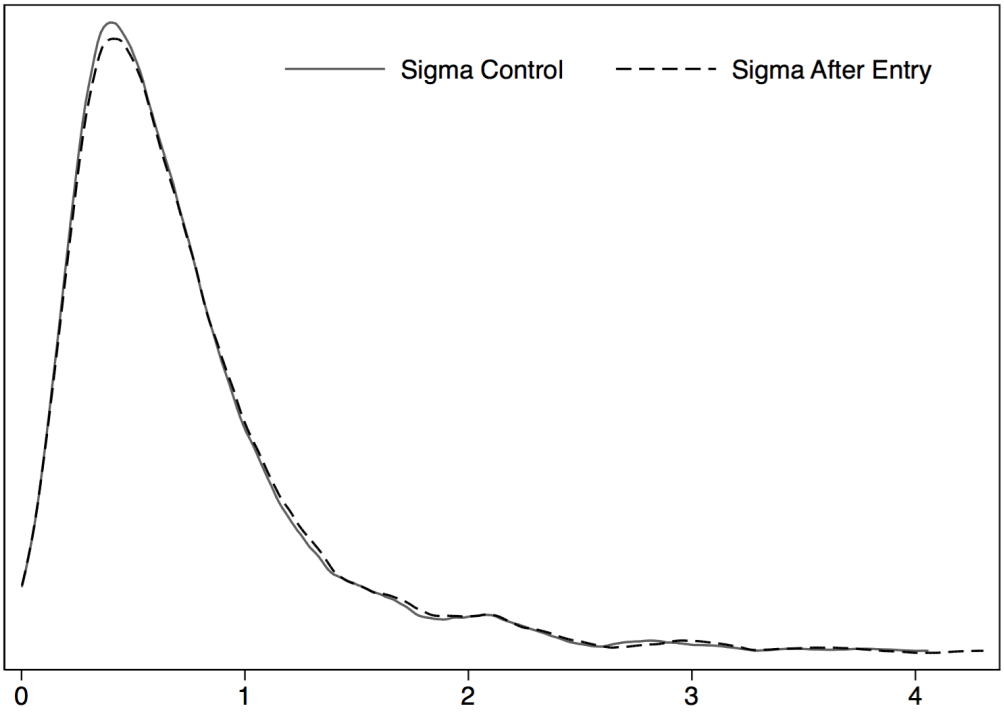

Figure 16: **Distribution of change in sigma**. The distribution of sigma T - sigma C from 1,000 bootstrapped samples is plotted. The mass is very close to 0 change (the outcome if the entrant colludes), with only slight increases in competition as seen by a slight shift to the right of the distribution. However, I cannot rule out no increase to competition, as the 90% confidence interval includes 0. The confidence interval does, however, exclude the predicted changes to sigma that one would expect to see if conduct (θ had remain unchained upon entry or if the entrant had competed (estimates of the change in sigma one would expect under each scenario represent the average change one would expect, given the distribution of the number of traders in each market in my sample).

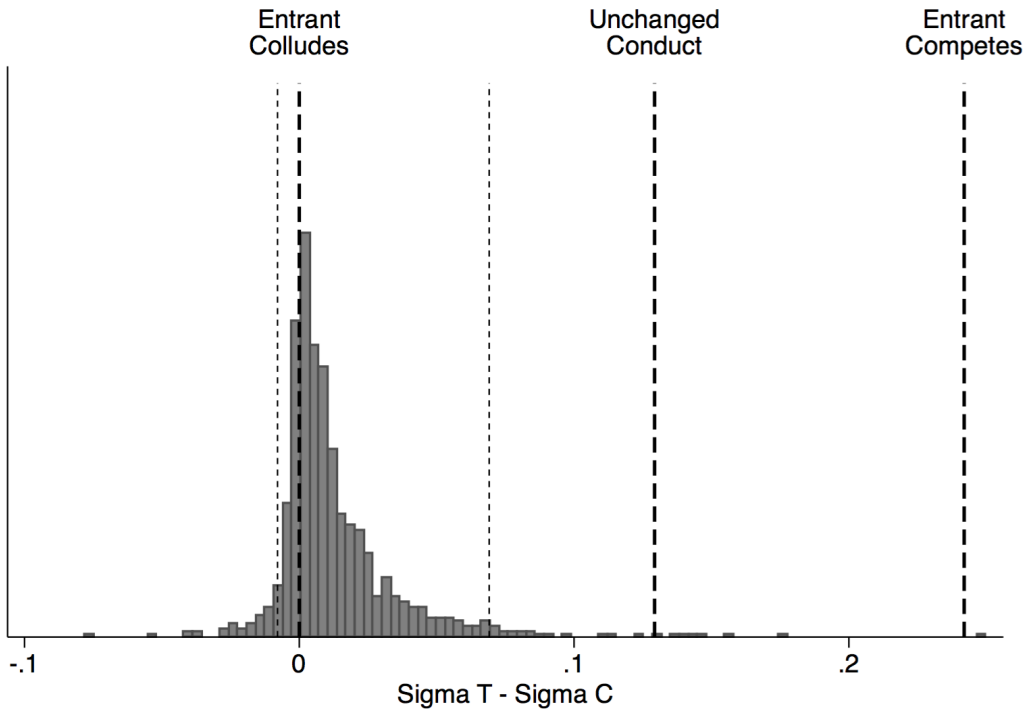

Table 1: **Trader summary statistics.** Data drawn from trader price surveys.

	Mean	Std. Dev.	Obs
Education and Business Characteristics			
Complete primary	0.78	0.42	2,728
Complete secondary	0.33	0.47	2,728
Percent corrrect Ravens	0.49	0.22	2,681
Review financial stregth monthly+	0.62	0.49	2,728
Keep written records	0.58	0.49	2,728
Any employees	0.37	0.48	2,728
Number employees	1.04	1.98	2,728
Own lorry	0.35	0.48	2,992
Market Experience			
Work in this market most weeks	0.95	0.22	2,964
New trader	0.02	0.13	2,964
Worked with all before	0.77	0.42	3,038
Know other traders well	0.67	0.47	2,571
Know other traders well or somewhat well	0.94	0.24	2,571
Collusion Reports			
Self-report discuss price	0.38	0.49	2,571
Someone in market report discuss price	0.80	0.40	2,806
Percent traders with whom discuss price	0.77	0.28	977
Self-report agree price	0.30	0.46	2,571
Someone in market report agree price	0.72	0.45	2,806
Percent traders with whom agree price	0.77	0.28	778

Table 2: **Pass-through.** The first column show the overall pass-through rate of 22%. The second column shows pass-through rates separately by "low" and "high" offers, which are remarkably similar.

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

Table 3: **Demand Estimation.** The point estimates and the 95% confidence intervals for the three estimated parameters of the Bulow-Pfleiderer demand function are displayed

	Parameter Estimate	**Lower Bound**	**Upper Bound**
a	42.50	42.22	57.60
b	0.0006	0.0000	0.0671
δ	2.80	0.01	5.89

Table 4: **Welfare Estimates.** The first row shows the point estimates on consumer surplus, intermediate surplus, and deadweight loss at the closest theory-consistent σ of 1 (and the correspondent theory-consistent pass-through rate of 26%). The second row presents the upper 95% confidence interval presents estimates of consumer welfare, which is maximized when calculating using the upper end of the confidence interval on ρ and σ.

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

Table 5: **Welfare Counterfactuals.** The first row shows the point estimates on consumer surplus, intermediate surplus, and deadweight loss at the closest theory-consistent σ of 1 (and the correspondent theory-consistent pass-through rate of 26%). The second row presents the counterfactual welfare distribution if markets were Cournot competitive for the average market of four traders. The third row presents the counterfactual welfare if markets were perfectly competitive.

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

Table 6: **Take-up of Entry Offers.** Offers ranged from 5,000-15,000 Kenyan shillings ($49-148 USD). Take-up rate = 1 if the trader *ever* took up an offer during any of the four weeks for which the offer was available. In a survey conducted during the offer phone call, traders reported the revenues and costs they would expect to incur if they did not take up the offer and instead followed their typical schedule and, similarly, the revenues and costs they would expect to incur if they did take up the offer. "Report profitable"=1 when the profits expected under take-up + the offer amount > profits expected under no take-up. Reported here are revenues and costs for the full week surrounding the offer, in order to capture any effect take-up might have on other sales in the week, if travel time or inventory effects drive inter-temporal trade-offs (this is the conservative choice; "reported profitable" rates are slightly higher if I use daily profits). Note that for only this column, the number of observations is 167 (rather than 180), as 13 traders could not be reached or refused to participate in the offer survey.

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

Table 7: **Take-up of Entry Offers by Market Size.** Outcome variable regressed on week fixed effects, market fixed effects, dummies for the number of traders present in the market (as categorized at baseline in Figure 2), and interactions of each of these dummies and an indicator for entry treatment. Only interaction term coefficients are displayed here.

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

Table 8: **Effect of Entry.** The variable "Entry Offer Market" is a dummy for treatment status in the entry experiment. "Num Traders" is the number of traders present in the market on that day. Column 1 presents the first stage effect of treatment on the number of traders. Column 2 presents the reduced form effect of treatment on log price. Column 3 presents the effect of the number of traders on the log price, instrumenting for the number of traders with treatment. All specifications include market and week fixed effects.

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

Table 9: **Simulated effect of entry on prices for various models of entrant behavior.** Estimated effect on θ, σ, and prices under various forms of entrant behavior. Effects are identified separately at each level of baseline number of traders in the market. Initial baseline σ_C is assumed to be 1 and demand parameters are taken at their point estimate. The **top panel** presents simulated IV effects of impact of entry by one trader. The **middle panel** presents simulated reduced form effects of the impact of entry, given the first-stage increase in the number of traders observed at each market-size. Because these first-stage effects may contain substantial noise, given the small cells of some market-size buckets, I present in the **bottom panel** simulated reduced form effects of the impact of entry using the average first-stage for all market-sizes.

Baseline	**Num**	**Δ N**	**Conduct Unchanged**			**Entrant Competes**			**Entrant Colludes**		
Num Traders	**Mkts**		*Theta*	*Sigma*	*% Price Δ*	*Theta*	*Sigma*	*% Price Δ*	*Theta*	*Sigma*	*% Price Δ*
IV											
1 Trader	5	1.00	1	2.00	0.27	1.00	2.00	0.27	1.00	1	0
2 Traders	12	1.00	2	1.50	0.15	1.67	1.80	0.23	0.50	1	0
3 Traders	14	1.00	3	1.33	0.10	2.50	1.60	0.18	0.33	1	0
4 Traders	8	1.00	4	1.25	0.08	3.40	1.47	0.14	0.25	1	0
5 Traders	5	1.00	5	1.20	0.06	4.33	1.38	0.12	0.20	1	0
6 Traders	3	1.00	6	1.17	0.05	5.29	1.32	0.10	0.17	1	0
7 Traders	8	1.00	7	1.14	0.05	6.25	1.28	0.09	0.14	1	0
8 Traders	1	1.00	8	1.12	0.04	7.22	1.25	0.08	0.12	1	0
9 Traders	3	1.00	9	1.11	0.04	8.20	1.22	0.07	0.11	1	0
10 Traders	1	1.00	10	1.10	0.03	9.18	1.20	0.06	0.10	1	0
By Market RF											
1 Trader	5	-0.02	1	0.98	-0.01	1.00	0.98	-0.01	1.00	1	0
2 Traders	12	0.28	2	1.14	0.05	1.88	1.22	0.07	0.50	1	0
3 Traders	14	0.29	3	1.10	0.03	2.82	1.17	0.05	0.33	1	0
4 Traders	8	0.80	4	1.20	0.06	3.50	1.37	0.11	0.25	1	0
5 Traders	5	1.08	5	1.22	0.07	4.29	1.42	0.13	0.20	1	0
6 Traders	3	0.91	6	1.15	0.05	5.34	1.29	0.09	0.17	1	0
7 Traders	8	0.20	7	1.03	0.01	6.83	1.06	0.02	0.14	1	0
8 Traders	1	1.76	8	1.22	0.07	6.74	1.45	0.14	0.12	1	0
9 Traders	3	1.08	9	1.12	0.04	8.14	1.24	0.08	0.11	1	0
10 Traders	1	8.43	10	1.84	0.24	5.88	3.13	0.47	0.10	1	0
Average RF											
1 Trader	5	0.58	1	1.58	0.17	1.00	1.58	0.17	1.00	1	0
2 Traders	12	0.58	2	1.29	0.09	1.77	1.45	0.14	0.50	1	0
3 Traders	14	0.58	3	1.19	0.06	2.68	1.34	0.11	0.33	1	0
4 Traders	8	0.58	4	1.15	0.05	3.62	1.27	0.08	0.25	1	0
5 Traders	5	0.58	5	1.12	0.04	4.58	1.22	0.07	0.20	1	0
6 Traders	3	0.58	6	1.10	0.03	5.56	1.18	0.06	0.17	1	0
7 Traders	8	0.58	7	1.08	0.03	6.54	1.16	0.05	0.14	1	0
8 Traders	1	0.58	8	1.07	0.02	7.53	1.14	0.05	0.12	1	0
9 Traders	3	0.58	9	1.06	0.02	8.51	1.13	0.04	0.11	1	0
10 Traders	1	0.58	10	1.06	0.02	9.51	1.11	0.04	0.10	1	0

A Appendix: Product Differentiation

I observe little variation in quality, as measured on a scale of 1-4 (97% of all maize receiving a rating of 2 or 3). Moreover, as shown in Column 1 of Table A.1, prices are not statistically different across the (limited) variation seen in quality. The other salient dimension on which products might vary is credit. However, credit does not appear to be a major factor in these primarily "cash-and-carry" spot markets; over 95% of transactions are conducted in cash.[78] Moreover, while I do see small price differences for purchases on credit, this relationship disappears when controlling for other features of the transaction.[79]

Therefore, the weight of evidence appears to suggest that maize sold in these markets is a relatively homogenous good. That said, I address the impact of the existing of differentiated products on my model predictions in Section 3. I show that a market with perfectly differentiated products generates identical predictions to that of perfectly collusive behavior in a market with homogenous products (under the assumption of similar curvature across demand for different product types). Therefore, while the changes the interpretation of the method by which market power is exerted – collusion vs. differentiated products – it does not alter the overall findings of a large degree of market power.

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

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

[78] That said, it may be that the *availability* of credit matters to a minority of customers. When asked how customers decide on which trader from whom to buy, 34% cite the availability of credit when needed, so it does appear that a slightly larger percent of customers value the possibility of obtaining a line of credit in periods when they are in need (results available upon request).

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

B Appendix: Price Discrimination

Recent work by Attanasio and Pastorino (2015) suggests that sellers of basic food staples in Mexico are able to exert market power to discriminate across customers with different levels of willingness (and ability) to pay. Using bulk discounts, sellers in their setting offer nonlinear pricing schemes in order to charge different prices to different consumer groups. Here I explore whether I see evidence of such nonlinear pricing schemes in my setting. To do so, I utilize transaction-level data (totaling 39,667 transactions) and explore the covariance of price and quantity of maize sold by the same trader to his customers in a given market-day. Figure B.1 presents this relationship, plotting a kernel-weighted local polynomial regression of log price on log quantity, both demeaned by trader x market-day fixed effects. While the relationship is relatively flat in the middle of the distribution, I see that customers at the lower end of the quantity distribution are paying more per kg, while those at the higher end are paying less per kg. The 95% confidence interval area, delineated in grey, suggests that these bulk discounts are particularly prominent at very large quantities. While the effect sizes are relatively small (the bulk of overall variation of price lies within a band of about 1%), they do suggest that traders possess some ability to use nonlinear pricing to price discriminate.[80]

Figure B.1: **Price discrimination and quantity discounts**. *Within trader x market-day residuals of transaction-level log price (per kg) and quantity (kg). N = 39,667. The grey area represents the 95% confidence interval area.*

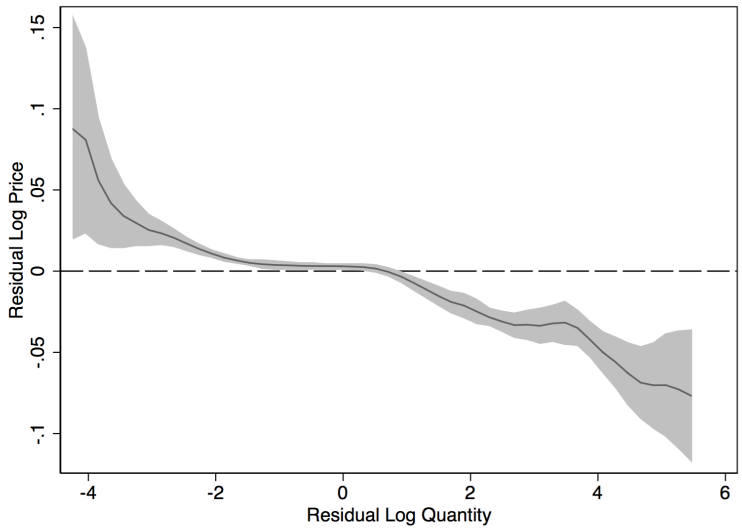

[80]Note that while incorporating the existence of some small degree of price discrimination would slightly alter the precise model predictions for pass-through under collusion, which rely on a specific trade-off between price charged and quantity sold, it would not affect the benchmark predictions for pass-through under perfect competition and Cournot competition (which by definition preclude price discrimination). Because I am unable to estimate individual demand curves, nor do I know the degree to which the trader can price discriminate, optimal pass-through rates under price discrimination are indeterminate (though ongoing work by the author in a parallel paper will attempt to identify some of these features broken down by groups salient to the trader, such as consumer type, bulk transactions, etc.). For now, I merely note that the ability to price discriminate at all is further evidence of market power, as traders in perfectly competitive or Cournot competitive markets are unable to price discriminate.

C Appendix: Sample Selection

The sample of markets in this study is drawn from six counties in Western Kenya. These counties encompass most of the (Kenyan) area within a 50km radius from the town of Bungoma, Kenya, the site of the research hub for this study. A listing exercise was conducted with the Director of Trade in each county to get a comprehensive list of all markets in the county. I excluded markets that were reported to typically not have any maize traders present. These represent some of the smallest rural markets, which have only maize retailers, who in turn purchase their maize from traders in larger markets. Major urban markets in the town centers were also excluded since the primary focus on of this study is on the rural markets frequented by rural consumers.[81]

The exercise yielded 154 potential markets for inclusion. From this sample, 60 markets were selected in the following stratified manner: 40 markets were selected from within a radius of 50 km of Bungoma town and 20 markets were selected from outside this radius.[82] I administered a pre-experiment survey of to this group of 60 selected markets in which I verified information provided by the Director of Trade and recorded the number of traders typically in the market.[83]. In a large number of these markets, it was found that the information provided by the Director of Trade was inaccurate.[84] Markets that were deemed ineligible upon visit were then replaced with market from their same stratum.[85] Newly selected markets were then visited in an identical verification exercise. This process was continued until 60 markets had been selected for inclusion in the sample.

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

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

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

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

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

D Appendix: How do Entrants Compare to Incumbents?

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

$$Y_{idw} = \alpha + \beta E_{idw} + \lambda_{dw} + \epsilon_{idw} \tag{24}$$

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

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

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

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

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

Table D.2: **Effect of Entry on Incumbents-Only.** *Column 1 presents the main IV specification from Table 8, while Column 2 presents the same specification with entrant traders removed from the sample.*

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

E Appendix: Duration and Intensity of Entry

It may be of interest to observe how variation in duration of entry – that is, the number of weeks in a row in which an entrant attends a market — and intensity – the number of entrants who attend a market at once – are correlated with price effects. One might expect, for example, that longer-run entry does more to increase competition (or the opposite: perhaps longer-run entrants are able to develop collusive relationships with incumbents through repeated interaction). Similarly, one might expect that the effect of entry is non-linear in the number of entrants: perhaps one entrant cannot do much to alter the market environment, but two entrants may have a large impact. I explore these questions here by documenting how observed variation in duration and intensity of entry is correlated with price effects. However, it should be noted (and it should be obvious) that take-up decisions are endogenous choices, and that these results can therefore in no way be seen as causal (and may be in fact so muddied by endogeneity as to be totally uninformative, but the author leaves that assessment to the reader).

E.1 Duration

To explore the evidence on entry duration, I construct a measure M_i for each potential entrant i. This variable which takes on values from 0 to 4, records the number of times the trader entered his or her offer market. For each market, I take the maximum of this variable over the three entrants assigned to that market:

$$\max\{M_1, M_2, M_3\}$$

Table E.1 presents the number of traders in each category. I see that while the model market had no entry, among the markets that had entry, the model market had three-weeks of entry and 66% of all markets with any entry experience 3 or 4 weeks of entry, suggesting that main entry observed in the experiment tends to be fairly long-run. Still, I do see some variation in the duration of entry, which I explore below.

Table E.1: **Duration of Entry.** *Each market is categorized by the maximum number of times entered across their three potential entrants.*

	Number Markets	**Percent**
No entry	28	0.47
One-time entry	4	0.07
Two-week repeat entry	7	0.12
Three-week repeat entry	13	0.22
Four-week repeat entry	8	0.13

I then run the standard 2SLS specification of the effect of entry separately for each category of market, identical to that outlined in Equation 18:

$$lnP_{iwd} = \alpha + \beta \hat{N_{wd}} + \gamma_w + \zeta_d + \epsilon_{it}$$

$$N_{wd} = \alpha + \beta T + \gamma_w + \zeta_d + \epsilon_{it}$$

Figure E.1 presents the point estimates and 95% confidence intervals on $\hat{N}$ from each of these separate IV regressions. While sensibly I see no effect in markets in which take-up was zero, and the negative effect seen seems concentrated among the markets that received 2+ weeks of entry, it is difficult to say more than that, given the size of the standard errors (and the broader concerns of endogeneity of this response).

Figure E.1: **Effect of Longer Duration Entry.** *Point estimates and 95% confidence intervals of a regression of log price on the number of traders, instrumented for with the entry treatment status. Regressions estimated separately for each category of market from Table E.1.*

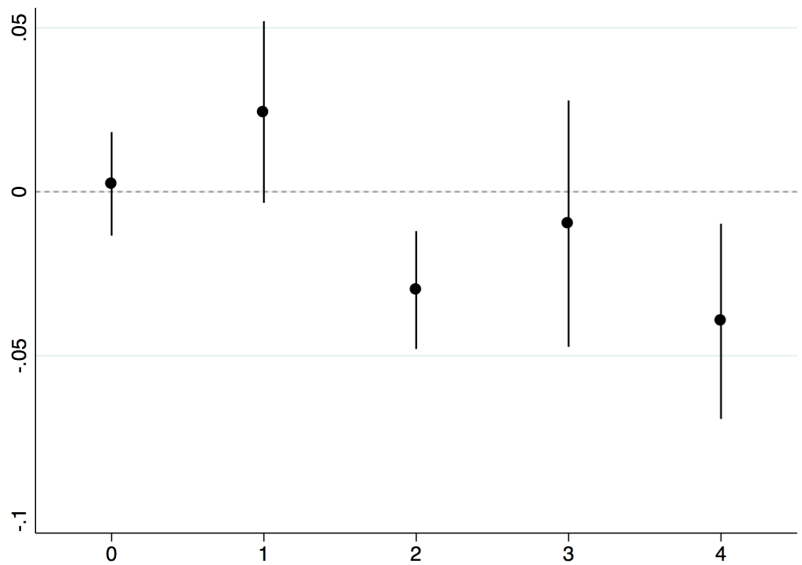

E.2 Intensity

To construct a measure of intensity of entry, for each market I average the number of entrants across the four treatment weeks:

$$\frac{1}{4}\sum_{w=1}^{4} Number\ Entrants_w$$

I then categorize markets based on whether they saw: no take-up, on average one entrant or less per market-day, or on average greater than one entrant. Note that there are very few markets in the greater than one category, and therefore this analysis will be grossly underpowered (as well as confounded by potential endogeneity).

Table E.2: **Intensity of Entry.** *Markets categorized by the average number of entrants seen in that market, averaged across the four market-days in which they were in the entry treatment status.*

	Number Markets	**Percent**
No take-up	28	0.47
Average entry of one or less	23	0.38
Average entry of more than one	9	0.15

I then run the standard 2SLS specification of the effect of entry separately for each category of market, identical to that outlined in Equation 18:

$$lnP_{iwd} = \alpha + \beta \hat{N_{wd}} + \gamma_w + \zeta_d + \epsilon_{it}$$

$$N_{wd} = \alpha + \beta T + \gamma_w + \zeta_d + \epsilon_{it}$$

Figure E.2 presents the point estimates and 95% confidence intervals on $\hat{N}$ from each of these separate IV regressions. While I sensibly again see effects concentrated in markets that experienced any entry, point estimates oddly suggest a *bigger* price decrease in markets that saw on average one trader or less compared to those that saw more than one. However, these point estimates are not measured precisely (point estimates overlap) and the small sample of markets with more than one entrant on average (as well as overall endogeneity concerns) make clear interpretation difficult.

Figure E.2: **Effect of Greater Intensity of Entry.** *Point estimates and 95% confidence intervals of a regression of log price on the number of traders, instrumented for with the entry treatment status. Regressions estimated separately for each category of market from Table E.2.*

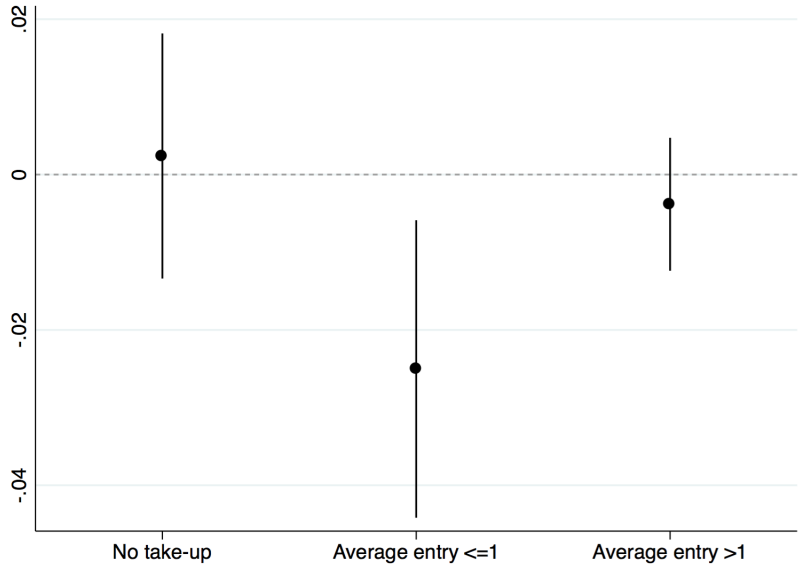

F Appendix: Long-Run Effects

It was necessary to rotate markets through the two treatment periods and one control period so that I could to include market and week fixed effects to soak up the substantial variation observed in prices across markets and over the season.[86] One additional benefit of this design is that it allows me to test for any long-run effects of the treatments, which were offered in a random order. Table F.1 presents these long-run effects.

First, for comparison's sake, I present in Column 1 the following specification:

$$lnP_{iwd} = \alpha + \beta_1 PT_{wd} + \beta_2 Ewd + \gamma_w + \zeta_d + \epsilon_{it}$$

in which lnP_{iwd} is the log price charged by trader i in week w in market d. $PTwd$ is a dummy for whether market d is in the pass-through experiment treatment during week w and E_{wd} is similarly a dummy for whether market d is in the entry experiment treatment during week w. As before, γ_w are week fixed effects and ζ_d are market fixed effects.

I present Column 1 with the full sample for comparison.[87] I then present the same specification, but with the sample restricted to only the second and third 4-week blocks in Column 2. I do this again for comparison purposes, as the spillover results in Column 3, which take into account the previous block's treatment status, will drop Block 1.[88] Column 3 runs an nearly identical specification to Column 2, but adds dummies for the market's treatment status in the previous block. It is the coefficients on these previous treatment status dummies that are of interest here.

I see no evidence of any long-run effects from the pass-through experiment, as the point estimate on "PT Previous" is small and statistically indistinguishable from zero. This should not be surprising, given the lack of price stickiness observed in these markets generally (see Figure 4). Clear communication of the duration of the subsidy, which resulted in well-understood start and end dates, likely also contributed to this clean effect.[89]

However, I do observe some interesting long-run effects from the entry experiment. This is particularly interesting given that no entrant returned following the cessation of treatment; any lingering effects, therefore, may be due to sustained changes in incumbent behavior.[90] The magnitude of the long-run effects are a little more than half the size of the main effect, a percent breakdown that is remarkably similar to that seen in Table D.2, in which I observed that incumbent reductions in price account for a little over half of the total reduction in price seen in the entry experiment. This further corroborates the hypothesis that the long-run changes observed may be driven by persistent incumbent behavior change, albeit quite small. The robustness of this result, as well as the potential mechanism driving this result, is an interesting area for further exploration.

[86] Piloting revealed that market and week fixed effects cut standard errors in half.

[87] This specification is different from the main specifications, because I am combining all experiments and including dummies for both treatment statuses. However, note that the coefficients line up in magnitude with the amounts estimated in the main specifications for the pass-through experiment and entry experiment. Therefore, Column 1 should be thought of as a benchmark for the results in the main specifications.

[88] Note that the treatment effects of both the pass-through experiment and the entry experiment are larger when I exclude block 1. It is not clear if this reflects heterogeneity in markets (perhaps something distinct about the 20 markets that served as controls during block 1) or in the seasons (perhaps something distinct about the early part of the season that occurred during block 1).

[89] The "buffer week" in between each block, during which time the demand experiment was run in a subset of markets, likely also mitigated any spillovers across treatments.

[90] These effects are unlikely to be due to price stickiness, since I see no spillovers from the pass-through experiment, which induced larger absolute price changes.

After restricting the sample and controlling for these long-run effects (which are essentially spillover effects, from the perspective of the overall study design), the reduced form effect on prices appears to be highly significant at 2.26%. This is substantially greater than the overall reduced form effect of 0.6% estimated in Column 1 (which is insignificant). However, it should be noted that a large part of this change is due to the restricted sample, as seen by the increase in the coefficient on "E" from Column 1 to Column 2.[91] Moreover, given a first stage effect of entry of about 0.61 traders in this restricted sample, this is consistent with an IV effect in which one trader entering is associated with a 3.7% reduction in price. While this is larger than that seen in the full experiment, this is still lower than the predicted IV effects delineated in the top panel of Table 9 under scenarios in which conduct is unchanged or the entrant competes upon entry (in which we would expect IV effects of 10% and 16% reductions in price, respectively) . Therefore, while it does suggest some small price decreases resulting from entry, I can still rule out effects consistent with the more competitive benchmarks.

Table F.1: **Long-run Effects** *Log price regressed on dummy indicators for pass-through (PT) and entry (E) treatment status. Column 1 presents the full sample, while Column 2 and 3 are restricted to blocks 2 and 3. Column 3 controls for the treatment status of the previous block. All specifications contain market and week fixed effects.*

	Ln Price	Ln Price	Ln Price
PT	-0.0207***	-0.0303***	-0.0317***
	(0.00447)	(0.00595)	(0.00619)
E	-0.00626	-0.0145***	-0.0226***
	(0.00418)	(0.00487)	(0.00622)
PT Previous			-0.00316
			(0.00513)
E Previous			-0.0157**
			(0.00609)
Mean DV	3.360	3.390	3.390
N	2802	2029	2029
Sample	Full	B2&3	B2&3
Market FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes

[91] Note that for the main results, there is no theoretical reason to restrict the sample in this way, nor was such a sample restriction pre-specified in the design registry. It is done here merely for practical purposes to look for long-run effects that could only possibly be seen in blocks 2 and 3.

G Appendix: Evolution of Effects

In this Appendix, I explore how effect sizes evolve over the four-week treatment period in both the pass-through and the entry experiment. Up front, it is important to note that the experiment spans the duration of the lean season, which runs from mid-March, when farmers in the region typically begin running out of their own stores of maize and begin turning to the market (and therefore prices begin to rise), to July, the peak of prices and the hunger season.[92] There are therefore likely underlying seasons trends over the course of the experiment, and any differences seen in effects seen across weeks should be interpreted with caution, as it is not clear whether these represent the more "long-run" effect of these interventions or are the result of underlying seasonal trends.

With that caveat in place, Table G.1 presents the evolution of the effects of the pass-through experiment and the entry experiment. Column 1 presents the evolution of the pass-through rate. The specification run is identical to that of Equation 10, but with interaction terms between the treatment CR and dummies for each week of the (4-week-long) block. I observe that the pass-through rate does appear higher in weeks 3 and 4 compared to weeks 1 and 2. It is unclear if this reflects changes to the structure of demand over the season or some adjustment to the cost shock on the part of traders.

Estimating seasonal changes to demand (and possibly to market structure itself) may be in interesting avenue for future research. Because the demand experiment was run at different points in the season, and because the precise timing of the demand experiment was randomized across markets, it may be possible to shed light on this by conducting the demand estimating separately for each block. However, such analysis with this sample is likely underpowered, as such analysis entails dividing the sample into quarters and, as shown earlier, the curvature term δ is already measured with substantial noise.

It is also possible, of course, that the increase in pass-through reflects some adjustment period. To address this, I can as a robustness check ignore weeks 1 and 2, during which time traders may have been acclimating to the cost shock, and use only the pass-through rate seen in weeks 3 and 4, when pass-through appears to plateau around 30%.[93] However, doing so does not change the overall conclusion of strong market power. Even the highest pass-through rate of 31.9% observed in week 3 is statistically indistinguishable from the collusive prediction of 26% pass-through, and is bounded away from the prediction of a Cournot competitive model of 55% pass-through with 99% confidence.

I similarly see an interesting evolution of the effects of the entry experiment. First, note that take-up of the offer evolves over the four weeks of treatment. It is not immediately clear why this would be the case (one simple possibility is that traders were reminded of the offer during the weekly follow-up surveys that were conducted and this encouraged greater take-up). I sensibly see that the reduced form effects on price therefore increase over the period as well. Of most interest is Column 4, which accounts for this increase in take-up by instrumenting for the number of traders. Even after accounting for differential take-up, I estimate a larger effect of each additional trader in later weeks. Though point estimates are insignificant in weeks 1-3, I do see a signifiant drop in prices of about 2% in week 4.[94] Comparing these effect to the expected IV effects in Panel 1 of

[92] Prices then drop sharply in August as harvesting begins.

[93] An F-tests suggests that one can reject equality of the effects with 95% confidence. However, one cannot reject the equivalence of weeks 1 and 2, nor the equivalence of weeks 3 and 4.

[94] Though an F-test suggests that one cannot reject the equality of the four effect estimates. Equality of the week 4 estimate and that of any of the other week also cannot be rejected.

Table 9 suggests that the observed effects – even at their maximum in week 4 – are still far from the expected IV effects under counterfactual scenarios of conduct remaining unchanged or the entrant competing (in which one would expect IV effects of 10% and 16% reductions in price, respectively). However, it does suggest some interesting evolution of effects. It is possible, for example, that entrants chose to stop colluding with incumbents in the final week, as collusive incentives break down in the final period. It is difficult to say definitely from the evidence here, but the dynamics of these effects are interesting avenues for future research.

Table G.1: **Evolution of Effects.** *Effects broken down by week (1-4) of the intervention. Column 1 shows the pass-through rate from the pass-through experiment by week. Columns 2-4 present the entry experiment effects by week, looking at first stage effect of treatment on the number of traders (Column 2), the reduced form effect of treatment on price (Column 3), and effect of the number of traders on price, instrumenting for the number of traders with treatment (Column 4).*

	Pass-Through	Num Traders	Ln Price	Ln Price
Week 1	0.131**	0.416*	0.00140	0.000694
	(0.0572)	(0.227)	(0.00618)	(0.0147)
Week 2	0.146**	0.495**	0.00635	0.0145
	(0.0570)	(0.190)	(0.00638)	(0.0272)
Week 3	0.319***	0.751***	-0.0115*	-0.0149
	(0.0596)	(0.213)	(0.00585)	(0.00969)
Week 4	0.291***	0.833***	-0.0175***	-0.0202*
	(0.0714)	(0.224)	(0.00617)	(0.0114)
Mean Dep Var	28.92	4.305	3.364	3.364
N	1860	2045	1776	1776
Reg Type	ρ	FS	RF	IV
Market FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes