Efficiency, Competition, and Welfare in African Agricultural Markets

by

Lauren Falcao Bergquist

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Edward Miguel, Chair
Professor Stefano DellaVigna
Professor Benjamin Faber
Professor Jeremy Magruder

Spring 2017

Efficiency, Competition, and Welfare in African Agricultural Markets

Abstract

Efficiency, Competition, and Welfare in African Agricultural Markets

by

Lauren Falcao Bergquist

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Edward Miguel, Chair

African agricultural markets are characterized by large variation in prices across regions and over the course of the season, suggesting poor market integration. This thesis explores the barriers that prevent various market actors from engaging in efficient arbitrage. Using experimental evidence and original survey data, I test for the existence of market failures that may limit integration and measure the efficacy of potential remedies to these market failures. In the first chapter, I quantify the degree of competition among the intermediates responsible for agricultural trade. In the second chapter, I explore whether entry by new intermediaries can enhance market competition. In the final chapter, Marshall Burke, Edward Miguel, and I test whether missing credit markets contribute to farmers' inability to arbitrage seasonal price fluctuations. Together, these three essays contribute to our understanding of agricultural market efficiency, competition, and barriers to arbitrage.

Each chapter employs experimental tests motivated by – and designed to speak directly to – economic theory. The first two chapters use randomized controlled trials to identify model parameters, while the third uses these trials to quantify general equilibrium effects and measure how such effects can shape the individual-level impacts of interventions. Methodologically, these essays exploit the clean causal identification generated by randomized controlled trials in new ways, in an attempt to shed light on the underlying organization of market institutions.

The first chapter of this thesis estimates the level of competition among intermediaries in Kenyan agricultural markets. There has long been concern that the wedge between the low price farmers receive for their produce and the high price consumers pay for their food – and the resulting loss in producer and consumer welfare – are driven in part by imperfect competition among the intermediaries that connect them. However, there has been little definitive evidence on the market structure in which these intermediates are acting. A lack of record-keeping on the part of traders precludes accounting assessments of profits. Further, identifying clean cost shocks and tracing pass-through is made difficult by the ubiquitous nature of production and consumption of agricultural commodities, which drives co-movement in supply and demand.

The first chapter overcomes these challenges by providing experimental estimates of pass-through and demand, key parameters governing the competitive environment of these markets. I identify these parameters using two randomized control trials that are tightly linked to a model of market competition. In the first experiment, I reduce the marginal costs of traders in randomly selected markets by offering to traders a subsidy per bag sold. I find that only 22% of this cost reduction is passed through to consumers. A second experiment offers randomized price discounts to consumers and measures corresponding quantities purchased in order to elicit the curvature of demand that traders face. I employ these estimates in a structural model of competition and optimal pricing to identify the level of competition among intermediaries. This exercise reveals a high degree of collusion among intermediaries, with large implied losses to consumer welfare and overall market efficiency.

The second chapter explores the impact of one natural policy response to this low level of competition: greater firm entry. In order to identify the impact of firm entry on competition, I randomly incentivize the entry of new traders into markets. I find limited benefit for consumers, as prices decrease only 1% in response to entry by one new trader. By capturing the resulting effect on local market prices, I identify the implied change in the competitive environment due to entry. This is most consistent with a model in which entrants are able to readily enter into collusive agreements with incumbents, suggesting that market power is robust to entry in this context.

The third chapter explores the barriers that limit arbitrage by farmers. Large and regular seasonal price fluctuations in local grain markets appear to offer African farmers substantial inter-temporal arbitrage opportunities, but these opportunities remain largely unexploited: small-scale farmers are commonly observed to "sell low and buy high" rather than the reverse. In a field experiment in Kenya, Marshall Burke, Edward Miguel, and I show that credit market imperfections limit farmers' abilities to move grain intertemporally, and that providing timely access to credit allows farmers to purchase at lower prices and sell at higher prices, increasing farm profits. To understand general equilibrium effects of these changes in behavior, we vary the density of loan offers across locations. We document significant effects of the credit intervention on seasonal price dispersion in local grain markets, and show that these general equilibrium effects strongly affect our individual level profitability estimates. In contrast to existing experimental work, our results indicate a setting in which microcredit can improve firm profitability, and suggest that general equilibrium effects can substantially shape estimates of microcredit's effectiveness.

Taken together, these results suggest that considerable inefficiencies exist in African agricultural markets. I find that agricultural traders in Kenya have considerable market power, and that marginal changes in market entry are unlikely to induce significant changes in competition. We further find that incomplete credit markets limit farmers' ability to arbitrage seasonal price fluctuations, and that the isolation of local markets reduces the sustainability of the financial products that may be necessary to encourage such arbitrage. These results have implications for the incidence of technological and infrastructure changes in African agriculture and for the policy responses aimed at improving the market environment.

To public higher education

Acknowledgments

I am deeply indebted to my dissertation committee members for their guidance and support. Working for Ted Miguel provided my first exposure to high-quality fieldwork, and his enthusiastic encouragement lead me to take risks in my own work. His advice and counsel, as well as his humor and kindness, were a constant source of support. I am extremely grateful to Ben Faber for his candor, guidance, and willingness to get into the weeds with me when I needed advice. The first two chapters of this dissertation would not have been possible without his generosity of intellect and time. Jeremy Magruder offered a constant sounding board, even for my craziest ideas, and I thank him for his patience. Our conversations were some of the first experiences I had of finding research to be fun. I thank Stefano DellaVigna for his unparalleled clarity as a teacher and mentor, and for showing me the value of linking models and empirics, which has profoundly shaped my research agenda.

I am grateful to Craig McIntosh, who treated me like a colleague before I deserved it. His feedback and guidance greatly improved this dissertation and his friendship is deeply appreciated. Ben Handel offered patient advice and constant support. I also thank Jeremy Weinstein and Michael Tomz, the originals in a chain of mentors who encouraged my foolish ambitions; they invested in me without basis, and changed the course of my life. Many faculty members at Berkeley gave helpful suggestions and generally enriched my graduate experience. These include Fred Finan, Supreet Kaur, Aprajit Mahajan, Andres Rodriguez-Clare, Betty Sadoulet and Reed Walker.

I am also grateful to Marshall Burke for his brilliance, humility and mentorship; to Willa Friedman, Sylvan Herskowitz, Joan Hamory Hicks, Kaushik Krishnan, Nicolas Li, Odyssia Ng, Jordan Ou, and Katalin Springel for their feedback and friendship; and to Jamie McCasland and Michael Walker in particular, whose advice and support kept me sane.

I thank Peter LeFrancois, Meshack Odhiambo Okello, Ben Wekesa, and Deborrah Muthoki Wambua for putting in endless hours managing these projects in the field, for being willing to try new things when they made sense and for reining me in when they did not, and for maintaining a sense of humor throughout it all. I also thank Innovations for Poverty Action and many skilled surveyors for excellent assistance in the field.

This research is supported by funding funding from the NSF Graduate Research Fellowship; the Agricultural Technology Adoption Initiative; an anonymous donor; 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).

Finally, I thank Jeff, for everything.

Contents

Chapter 1

Pass-through and Competition

1.1 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, Barret and Minten, 2009; Burke, Bergquist and Miguel, 2016).

High transaction costs contribute to this limited market integration. Transport costs in Africa are the highest in the world (Teravaninthorn and Raballand, 2009); also prevalent are harder-to-measure costs associated with search (Aker, 2010), contractual risk (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 two 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 and the shape of demand 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.

This paper is one of the first to experimentally test the competitiveness of rural agricultural markets directly. Previous attempts to measure competition have mainly relied on observational methods. Observational studies have typically found high rates of pass-through across major markets (Rashid and Minot, 2010), though these high transmission rates may not extend beyond major urban markets (Moser, Barret and Minten, 2009; 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. 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 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).[1] 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. 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).[2]

Finally, a set of papers has applied experimental methods to the somewhat related question of the impact of offering price information to farmers on their ability to extract better prices from traders. While most studies find null results (Fafchamps and Minten, 2012; Mitra et al., 2015),[3] it is unclear if this suggests traders are already offering competitive prices given their transport costs or whether farmers are simply unable to utilize this information to improve their bargaining position.[4] 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, Mookherjee and Tsumagari, 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

[1]The Kenyan 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.

[2]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.

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

[4]In a quasi-experimental variant of this literature, Casaburi, Glennerster and Suri (2013) 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.

setting in which ubiquitous production and consumption make it difficult to cleanly trace price shocks from distinct points of origin. In Chapter 2, I extend this work by identifying how key model parameters governing competition respond to entry.

This paper proceeds as follows: Section 3.2 describes the maize market industry in Kenya. Section 2.2 describes the theoretical model underpinning the experimental design, which is described in greater detail in Section 2.3. Section 1.5 presents results on pass-through, and Section 1.6 describes the demand estimation procedure. Section 1.7 presents the structural estimates of the level of competition among intermediaries and the welfare implications of these findings. Section 2.8 concludes and discusses policy implications.

1.2 Maize Markets in Kenya

Maize Output Market Chain

Figure 1.1 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. Regional traders, the subjects of this study, are responsible for large-scale aggregation, storage, and transportation. They report purchasing 50% of their maize from small and medium farmers (selling less than 5 tons), 16% from large farmers, and 33% from other traders. 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.

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. 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.

However, despite the lack of legal barriers, engaging in large-scale wholesale trade still requires significant working capital in order to pay for inventories, storage facilities,[5] and transport vehicles. 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). Further, traders must develop extensive networks of contacts in order to glean information on prices and product availability, as this information is disseminated one-on-one through personal networks of fellow traders rather than through any centralized or open information clearinghouse. It is common for traders to enter the business with the support of siblings, spouses, or even former employers who already have experience in the business. Therefore, while entry is close to free legally, those who wish to enter regional trade still face significant barriers.

Table 1.1 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.

Open Air Markets

This study takes place in the open air markets in which traders sell the majority of their produce. These markets typically occur on a set day each week. 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 1.2 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 1.1). 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%

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

"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.

Customers in these markets are comprised of two-thirds individual household consumers and one-third rural retailers. 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.[6] The product itself is fairly homogenous.[7]

1.3 Theoretical Framework

Chapter 1 explores the results of two 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. The experimental design is therefore tightly tied to theory. This section reviews that theory.

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:

$$\pi_{dt} = (P_{dt} - c_{dt})q_{dt} \tag{1.1}$$

Here, I employ a few simplifying assumptions. First, I assume that maize is a homogenous good and that traders are unable to price discriminate.There is little variation in quality and credit is rarely used. Further empirical evidence on product homogeneity is provided in Appendix A.1. 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

[6]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.

[7]See Appendix A.1 for additional evidence on the extent of product differentiation.

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).[8]

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, Gabre-Madhin and Minten (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 1.1 with respect to the trader's quantity q_{dt} yields the trader's first order condition:

$$P_{dt} = c_{dt} - \theta \frac{\partial P_{dt}}{\partial Q_{dt}} \frac{Q_{dt}}{N_{dt}} \tag{1.2}$$

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:

$$\theta = \begin{cases} 0 & \text{when perfectly competitive} \\ 1 & \text{when Cournot competitive} \\ N & \text{when perfectly collusive} \end{cases} \tag{1.3}$$

It is worth noting from Equation 2.1 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:"

$$\sigma \equiv \frac{N}{\theta} \tag{1.4}$$

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

$$\sigma \equiv \frac{N}{\theta} = \begin{cases} \infty & \text{when perfectly competitive} \\ N & \text{when cournot competitive} \\ 1 & \text{when collusive} \end{cases} \tag{1.5}$$

[8] 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 A.2, 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.

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 this chapter, which measures the competitiveness of these markets.

Pass-Through and Demand Curvature

To identify how traders respond to reductions in their marginal costs, taking the derivative of Equation 2.1 with respect to c_{dt} yields:

$$\rho_{dt} \equiv \frac{\partial P_{dt}}{\partial c_{dt}} = \left\{ 1 + \frac{1 + E_{dt}}{\sigma_{dt}} \right\}^{-1} \tag{1.6}$$

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.

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.1), I interpret market power as arising from collusion; nonetheless, they are observationally equivalent.

Figure 1.3 provides a visual example of the relationship between the market structure, the curvature demand, and pass-through. 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.

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 1.6

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):

$$\frac{IS}{CS} = \frac{1}{\bar{\rho}} + \frac{1-\sigma}{\sigma} \tag{1.7}$$

$$\frac{DWL}{IS} = (1-\bar{\rho}) + \bar{\rho}\sigma - \left(\frac{\bar{\rho}\sigma}{(1-\bar{\rho}) + \bar{\rho}\sigma}\right)^{\frac{\bar{\rho}}{1-\bar{\rho}}} (\bar{\rho}\sigma + 1) \tag{1.8}$$

where $\bar{\rho}$ is the quantity-weighted average pass-through rate.[9]

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 1.7 and 1.8 also allow for counterfactual simulations in which I evaluate the welfare implications of increases in the σ, the degree of competition.

1.4 Experimental Design

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.[10]

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 1.4) to minimize concerns about sticky prices.[11]

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.

[9] In Section 1.6, 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.

[10] See Appendix A.3 for additional details on the sample selection procedure.

[11] Appendix A.5 presents the effects of each treatment, broken down by week.

Figure 1.5 shows the six possible orders.[12], 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.

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 A.3 for further details on this census.

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. Only maize sold in cash was eligible for the subsidy due to concerns about the ability of enumerators to verify the authenticity of credit sales. However, over 95% of sales are conducted in cash, so this restriction was often irrelevant. The subsidy was capped at the first 75 90kg bags sold to limit budget exposure, but this cap was binding for only 1.5% of traders. When introducing the subsidy, enumerators first asked the trader to describe some of the major costs that he faced in his business.[13] 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

[12]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 A.4. Accounting for these long-run effects does not alter the overall implications of the results presented Section 2.4.

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

markets received a "high" subsidy level of 400kg/90kg bag. Sales of partial bags were eligible at the same prorated amount. Note that "low" and "high" are merely relative titles: both represent large and meaningful changes to traders' costs. The "low" subsidy rate represents 7.5% of the average price, while the "high" subsidy represents 15% of the average price. Payments were made via M-Pesa twice a day in order to build among traders.

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

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 in light of this new price. The price discount was given to the customer in the form of a mobile money or a cash transfer (consumer's choice) 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.

Experiment 3: Entry Experiment

A third "Entry" experiment was run with this sample over the same time period (although never run concurrently in single market; the order of the experiments was randomized across markets and therefore description of the randomization schedule requires mention of all three experiments. The content of the "Entry" experiment will be described in greater detail in Chapter 2.

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

Maize quality data was also collected for each trader. Enumerators were trained to grade quality on an objective scale from 1 (lowest quality) to 4 (highest quality). This was done because 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.

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.[14] Finally, the first time a trader was met in the sample, additional

[14]Note that, due to the sensitivity of these questions, these questions were asked at mid-day, after the enumerator had established good rapport with the respondent. However, there were several traders who

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

My primary outcome of interest – price – is defined as the quantity-weighted average of transaction level prices that the trader sold that day (because the ICC of price within a trader in a given market-day is high (0.9), in practice there is little variation in the prices entering into this average). When non-traditional reductions in price were offered, the value of these additional services were captured and incorporated into the price per kg.,[15]:

$$P_{idw} = \frac{\sum_{t=1}^{T} p_{idwt} q_{idwt}}{\sum_{t=1}^{T} q_{idwt}} \tag{1.9}$$

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. 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.

1.5 Pass-Through

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

Estimation

To measure pass-through, I estimate:

$$P_{idw} = \alpha + \beta CR_{dw} + \gamma_w + \zeta_d + \epsilon_{idw} \tag{1.10}$$

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.

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.

[15]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.

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

$$P_{idw} = \alpha + \beta_1 CR_{dw} * Low_{dw} + \beta_2 CR_{dw} * High_{dw} + \gamma_w + \zeta_d + \epsilon_{idw} \tag{1.11}$$

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 1.11, conditioning on the desired dimension of heterogeneity.

Results

Table 1.2 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. 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. The number of traders is defined as the average number of traders observed in the market over the course of the experiment. In order to remove any increases in the number of traders driven by the entry experiment, this figure uses the average of the predicted number of traders each week, based on market and week fixed effects. Figure 1.6 presents these results, which show little evidence of meaningful heterogeneity. Estimates of pass-through rates are fairly tightly centered around the overall estimate of 22% and no clear pattern is seen with the number of traders. To gain statistical power, the bottom two measures show the sample pooled into below and above median number of traders; again, point estimates are not statistically significantly different and are in fact remarkably close in magnitude.

In Figure 1.7, I further explore other sensible – although not pre-specified – dimensions of heterogeneity by a few other measures. 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.[16] The point estimates suggest that pass-through is sensibly smaller for markets above the median in this

[16] 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.

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.

1.6 Demand Estimation

As described in Section 2.2, 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.

Constant Pass-through Demand Class

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

$$Q_{dt}(P_{dt}) = \begin{cases} \left(\frac{a-P_{dt}}{b}\right)^{\frac{1}{\delta}} & if \ \ (P_{dt} \leq a, b > 0 \ \ and \ \ \delta > 0) \ \ or \ \ (P_{dt} \geq a, b < 0 \ \ and \ \ \delta < 0) \\ 0 & if \ \ P_{dt} > a, b > 0 \ \ and \ \ \delta > 0 \\ \infty & if \ \ P_{dt} \leq a, b < 0 \ \ and \ \ \delta < 0 \end{cases} \tag{1.12}$$

$$where \ a \geq 0$$

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 2.2, 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:

$$P_{dt} = a - bQ_{dt}^{\delta} \tag{1.13}$$

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 1.6 simplifies to:

$$\rho_{dt} \equiv \frac{\partial P_{dt}}{\partial c_{dt}} = \left(\frac{\sigma}{\sigma + \delta} \right) \tag{1.14}$$

Equation 1.14 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 1.2, 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.

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.[17] I estimate the vector of parameters $\Theta = (a, b, \delta)'$ in Equation 1.12 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 1.12 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:

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}}\, m(\Theta)'Wm(\Theta), \tag{1.15}$$

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.[18] 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 constrained by the results by the pass-through experiment).[19]

[17]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.

[18]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 $1.4x10^4$).

[19]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 1.3 for graphical intuition on this point).

Estimates are initialized at 500 randomly selected starting values, to ensure the minimization procedure does not obtain parameters for a local minimum. 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. 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 1.3, which show the point estimate and 95% confidence interval. Figure 1.8 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.

1.7 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. Given the point estimate on δ of 2.8, I use Equation 1.14 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: $\rho = \frac{N}{N+2.8}$. 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 1.9 displays the bootstrapped distribution of ρ. 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.

I see that the mass of the distribution of ρ is concentrated near the predicted pass-through of 26% under collusion. 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. Figure 1.10 presents this distribution, overlaid with the predicted σ under each model of competition. 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.

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. 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.

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 bounds and I am therefore able to reject them.[20]

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. As described in Section 2.2, perfect product discrimination would result in a σ of 1 as well. While the weight of the evidence presented in Appendix A.1 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.

Welfare Implications

What does this imply for the division of surplus between consumers and intermediaries? I use Equations 1.7 and 1.8 to solve for the ratios for consumer surplus (CS), intermediary

[20]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.

surplus (IS), and deadweight loss (DWL).[21] Table 1.4 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 1.7 and 1.8 the value of σ that corresponds to counterfactual forms of market conduct and the ρ that would be realized at each of these values of σ.[22] Table 1.5 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%.[23] 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 1.11 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.

1.8 Conclusion

Policymakers have long speculated that agricultural traders in Africa exert market power, paying below-competitive prices to farmers and charging above-competitive prices to consumers. However, the absence of trader records and the difficulty in identifying clean shocks to 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.

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

[22] As estimated by Equation 1.14, using the counterfactual σ and estimates of $\hat{\delta}$.

[23] 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 2.2, 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.

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.[24]

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 entry may be more a feasible response in this setting. It is the impact of these sorts of policies that I evaluate in Chapter 2.

[24]Some policies to reduce traders' costs may indirectly enhance competition. For example, road construction could encourage entry of new traders into markets.

Tables and Figures

Figure 1.1: **M**aize 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.

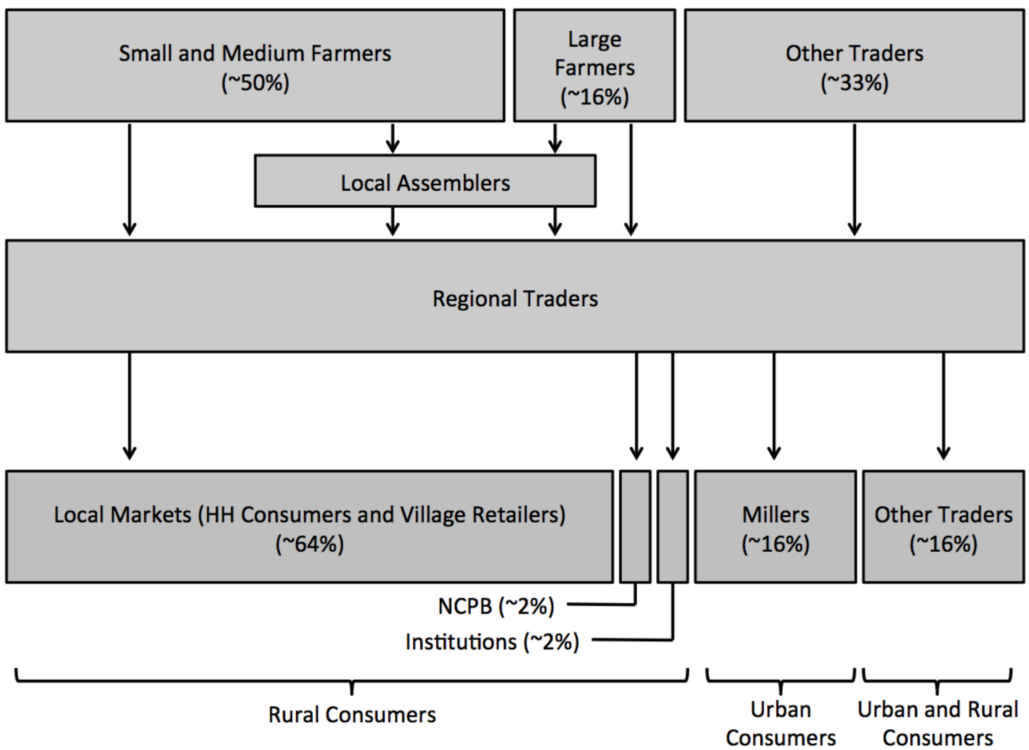

Figure 1.2: **H**istogram of number of traders. The number of traders is calculated as the average number of traders present in the market during 12 weeks of the study period, as predicted by week and market fixed effects (that is, any increase in number of traders due to the entry experiment is omitted).

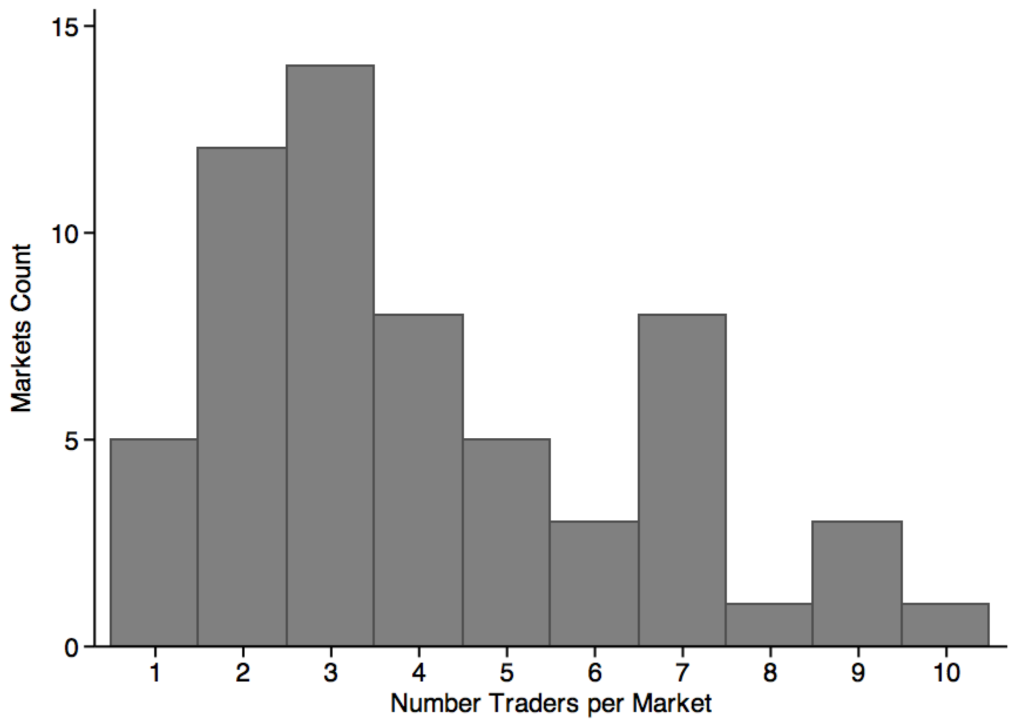

Figure 1.3: **P**ass-through given conduct and demand curvature. **L**eft panel: Here I show a collusive environment, in which traders act as a single profit-maximizing firm. Quantities are set where the marginal revenue (MR) curve meets the marginal cost (MC) curve. Prices are then set where this optimal 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. **R**ight 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.

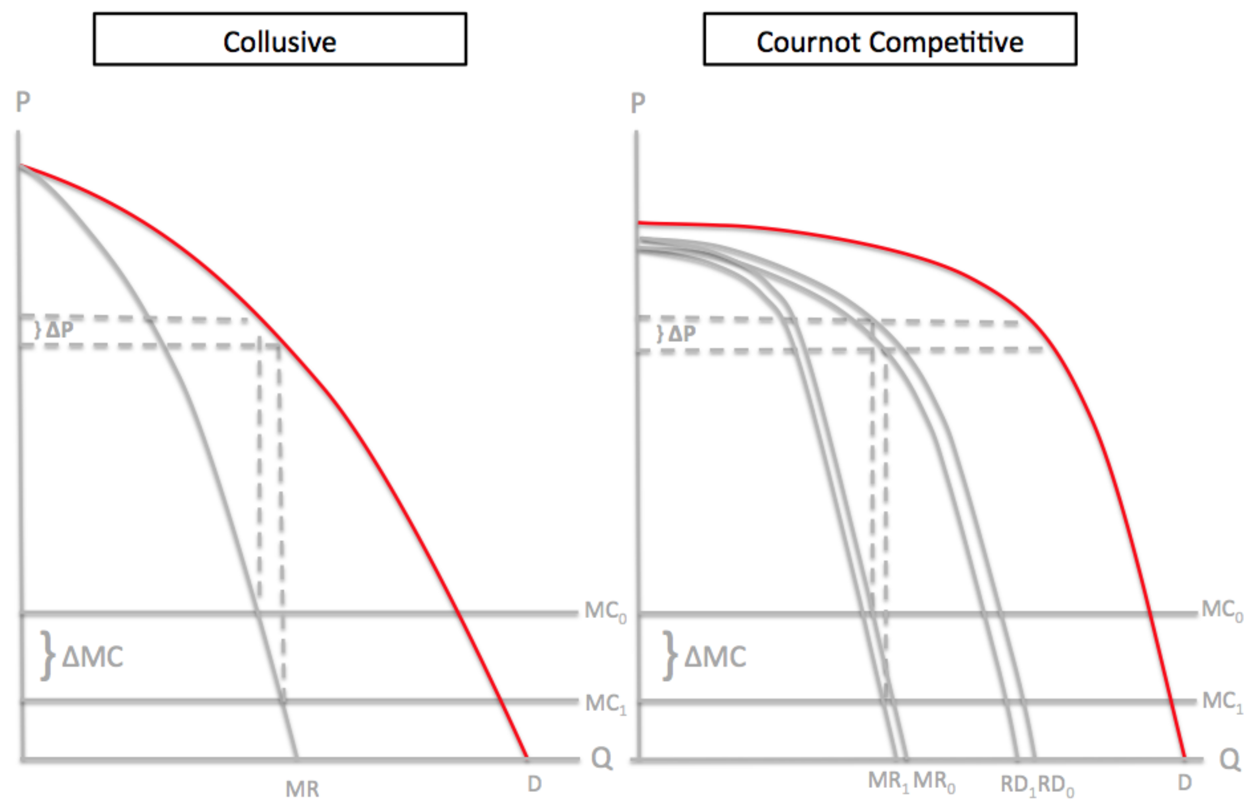

Figure 1.4: **M**aize prices in study markets. Grey lines show the price for each market over the 12 weeks of the study period. The black line shows the average price across markets in each week. The black bar on the vertical axis show the average size of the cost reduction subsidy for comparison. It 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.

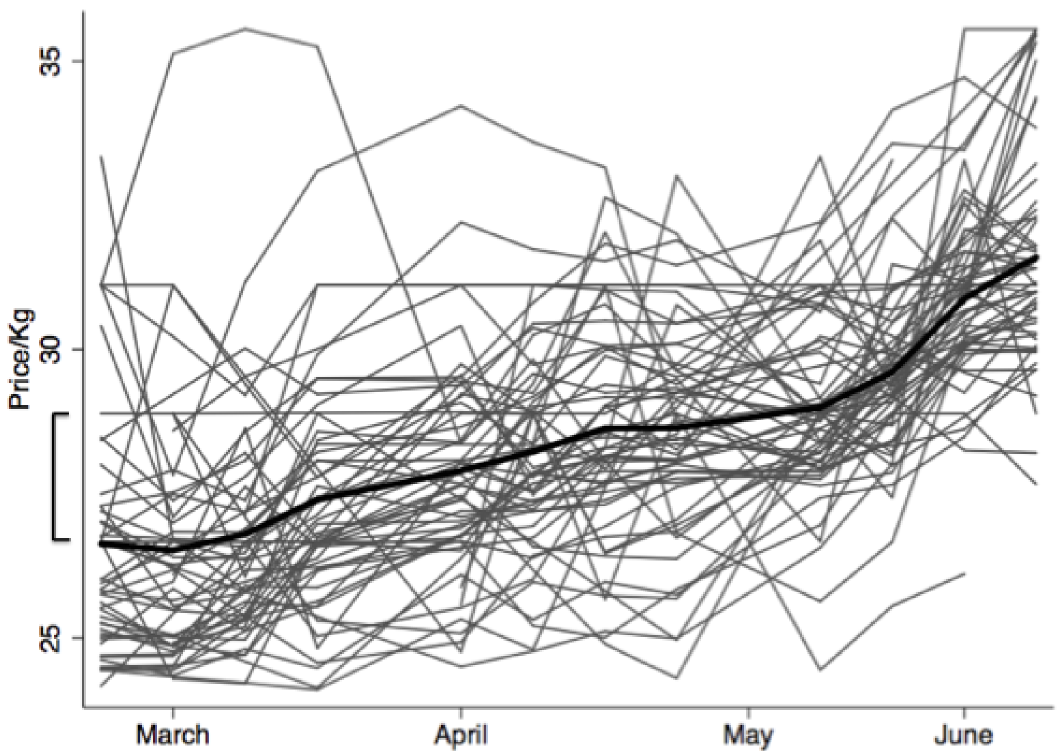

Figure 1.5: **E**xperimental schedule. The 60 markets in my sample are randomly assigned one of six possible schedules, in order to yield randomized ordering of treatment statuses. There are therefore 10 markets in each schedule. This allows the inclusion of market and week fixed effects in every analysis. 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.

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

Figure 1.6: **H**eterogeneity 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.

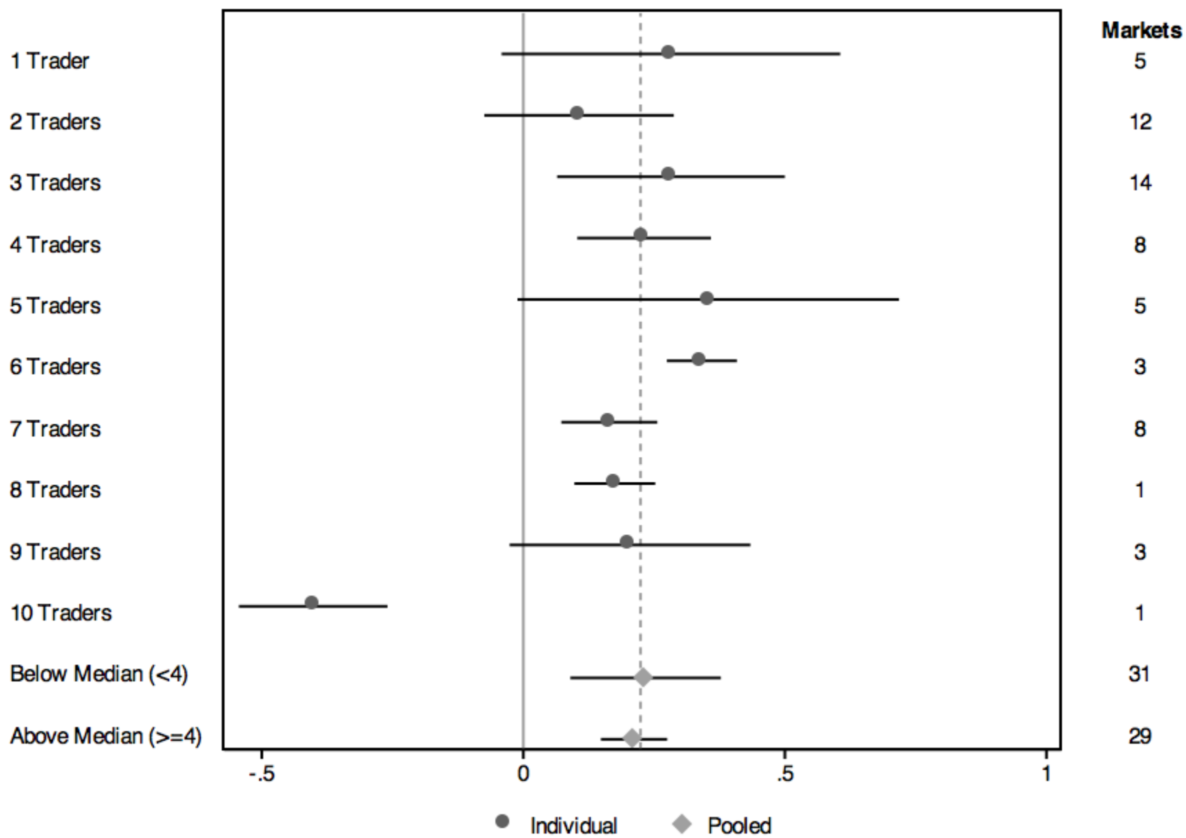

Figure 1.7: **H**eterogeneity in pass-through by various factors. No clear pattern of heterogeneity is seen in markets that are on tarmac roads (less geographically isolated) vs. off (more geographically isolated). The "discuss (agree) price" variable is constructed by counting the number of days in 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.

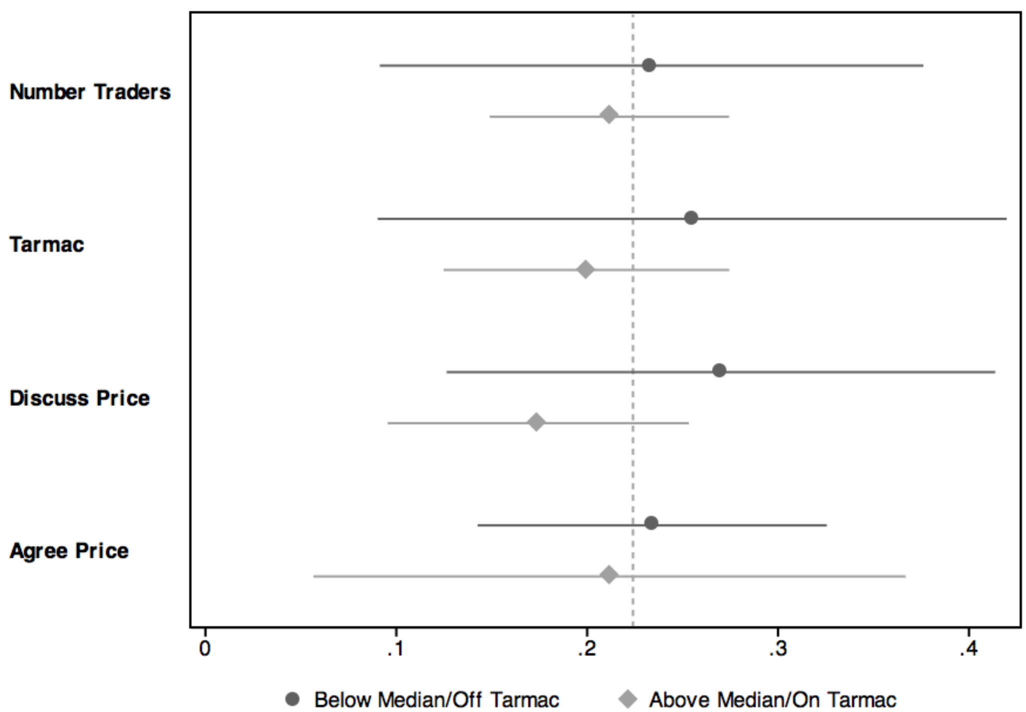

Figure 1.8: **D**emand fit. The x-axis plots the price as predicted by the (ten levels of) discounts offered in the demand experiment, while the y-axis plots the quantity purchased at that predicted price. The black line plots estimate Bulow-Pfleiderer demand functional form.

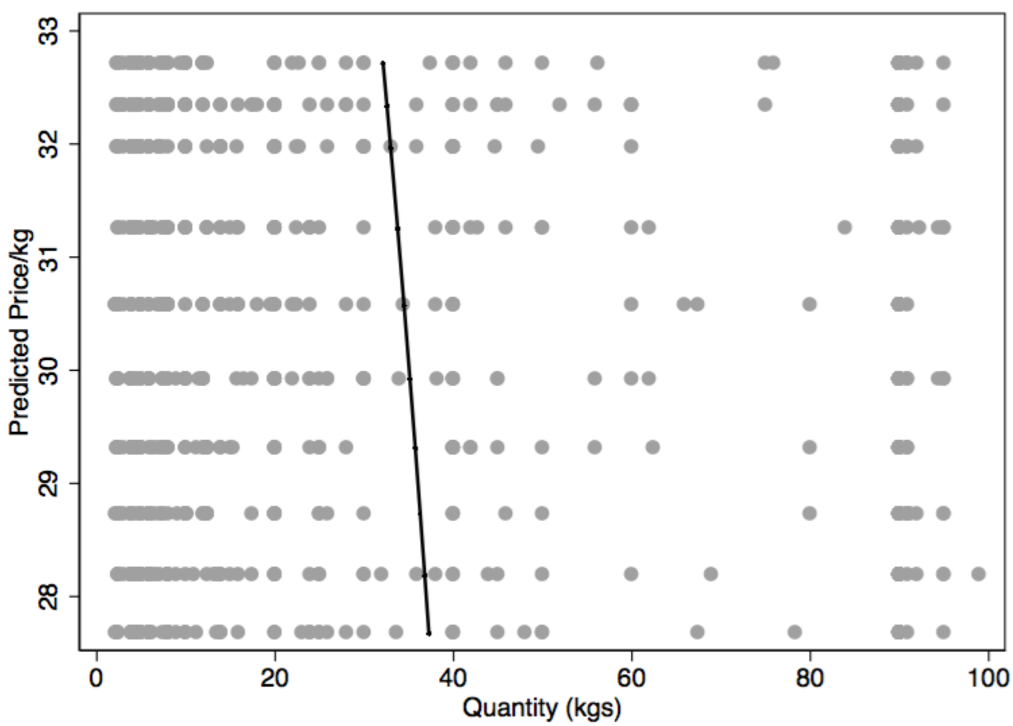

Figure 1.9: **P**redicted pass-through under various competitive environments. Given my estimate for the curvature of demand, I predict that one would have observed 100% pass-through in a perfectly competitive market, 55% pass-through in a Cournot competitive market, and 26% pass-through in a collusive 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.

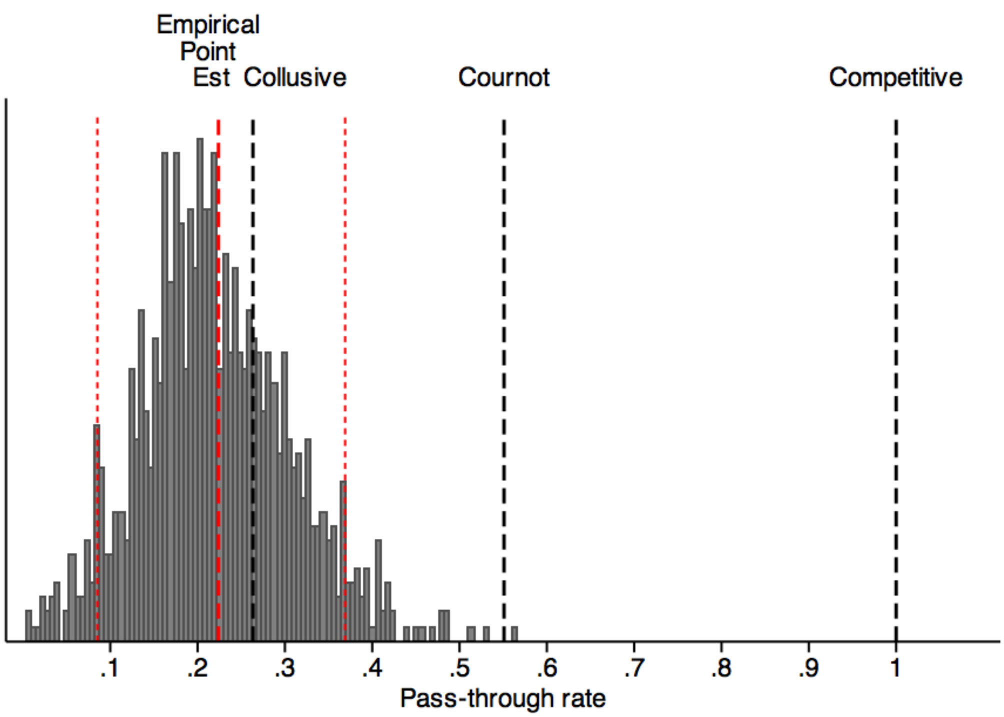

Figure 1.10: **C**ompetitiveness parameter (sigma) estimates. I combine my bootstrapped estimates of pass-through with my bootstrapped estimates of demand curvature in Equation 1.14 to generate 1,000 bootstrapped estimates of the competitiveness parameter σ. Recall that $\sigma = 1$ if competitive, N (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.

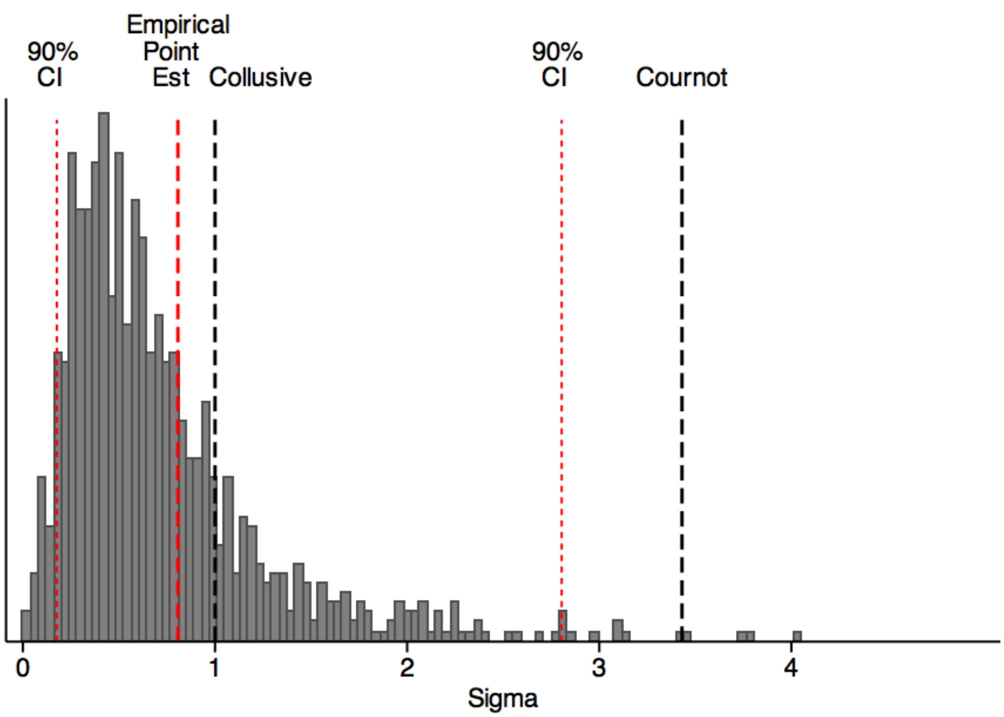

Figure 1.11: **W**elfare counterfactuals. Counterfactual division of welfare is shown for the average market of four traders. The current division of surplus is shown at the far left vertical dotted line, suggesting that intermediary surplus (IS) is 67.6% of total surplus, while consumer surplus (CS) 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.

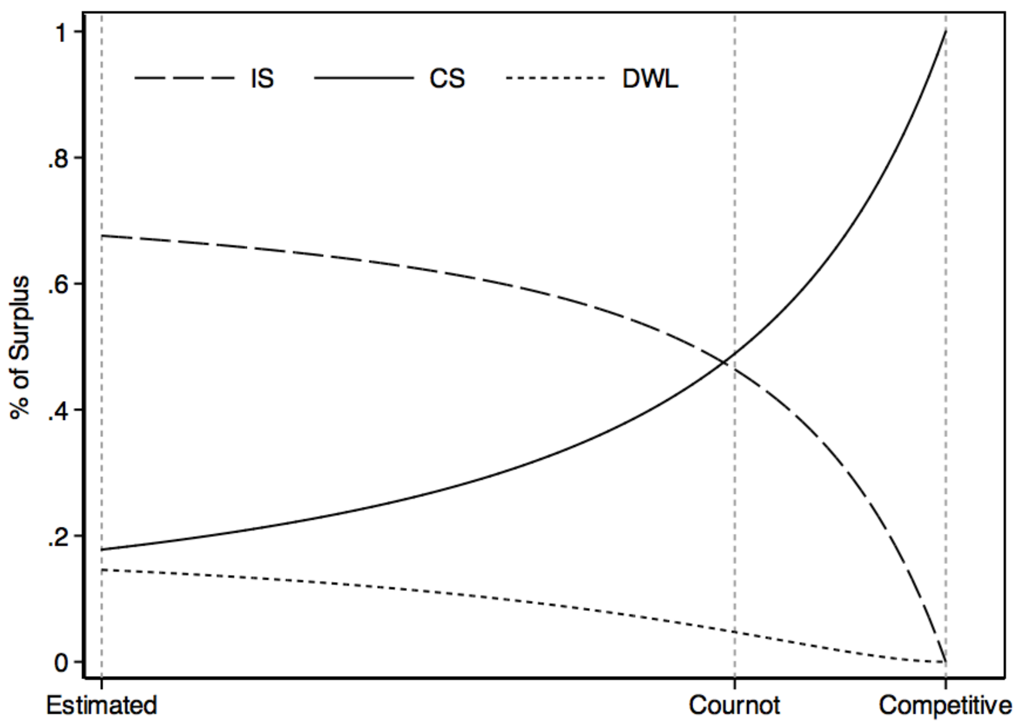

Table 1.1: **T**rader summary statistics. Data drawn from trader price surveys.

	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 1.2: **P**ass-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 1.3: **D**emand 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 1.4: **W**elfare 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 1.5: **W**elfare 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

Chapter 2

The Effects of Entry on Competition

2.1 Introduction

Given the finding from Chapter 1 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 the entry-inducing policies will enhance competition and improve consumer welfare. This is partially because identifying the empirical impact of entry has historically been difficult, as entry is typically an endogenous response to market conditions.

To overcome this challenge, I conduct an experiment in which I incentivize traders to randomly selected markets for the first time. These incentives drive an additional 0.6 traders per market-day on average, a 13% increase over the mean market size (and 20% over the median). I then use the exogenous entry generated by this experiment to test whether policies that incentivize market entry can decrease market power and promote competition.

In order to identify these effects, I use the model developed in Chapter 2 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.

2.2 Theoretical Framework

The theory employed in this experiment relies heavily on the framework developed in Chapter 1. In this experiment, the number of traders in the market is experimentally manipulated, and the effect of entry on both conduct and overall competitiveness is estimated.

Recall the trader's first order condition from Chapter 1:

$$P_{dt} = c_{dt} - \theta \frac{\partial P_{dt}}{\partial Q_{dt}} \frac{Q_{dt}}{N_{dt}} \tag{2.1}$$

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:

$$\theta = \begin{cases} 0 & \text{when perfectly competitive} \\ 1 & \text{when Cournot competitive} \\ N & \text{when perfectly collusive} \end{cases} \tag{2.2}$$

Recall we summarize the number of traders N_{dt} and how those traders interact (described by θ) with a single "competitiveness parameter:"

$$\sigma \equiv \frac{N}{\theta} \tag{2.3}$$

How will the level of competition σ change with entry? It is clear from Equation 2.3 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 this experiment.

2.3 Experimental Design

The sample for the entry experiment is identical to that employed for the two experiments described in the previous chapter. The Experimental Design section of the previous chapter describes the randomization schedule, as the three experiments are offered in a random order. It also describes the data collection process, which was identical as well. In this section, I describe in greater detail the 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 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. 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. 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.[1]

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. 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. 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.

[1]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.

2.4 The Cost of Entry

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 2.1 presents take-up at each subsidy level (take-up defined as ever accepting any of the four market day offers). Sensibly, I see that take-up increases in the size of the subsidy: take-up is 12% for the low offer, 28% for the medium offer, and 42% for the high offer. However, 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 (according to self-reports from a survey conducted on the phone at the time of the offer).

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.[2]

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:

$$\begin{aligned} T_{id} &= \alpha + \beta X_{id} + \epsilon_{id} \quad (i) \\ T_{id} &= \alpha + \beta X_{id} + \zeta_d + \epsilon_{id} \quad (ii) \end{aligned} \tag{2.4}$$

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 2.1 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 a drop of $46 USD in the offer amount.[3] Having contacts in the entry market is correlated with higher take-up (albeit not quite significantly). The value of having any

[2] 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.

[3] 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.

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. This is perhaps surprising, given work showing that ethnicity can matter for cooperation and productivity among employees in Kenya (Hjort, 2014). However, Hjort (2014) also finds that when pay is structured so as to reward the entire team based on joint performance, ethnic bias is diminished. It may be that in the setting of maize markets, in which a single market price influences the profits of all firms equally, it is more difficult to exert co-ethnic preference, and therefore ethnicity may not affect profits or in turn take-up of the entry offer. Further, the results are are also in line with experiment evidence showing a lack of co-ethnic bias in economic lab games (Berge et al., 2016).

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.

Appendix B.1 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. I turn now to the effect of this entry on prices.

2.5 The Effect of Entry on Price

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

$$logP_{idw} = \alpha + \beta EOM_{dw} + \gamma_w + \zeta_d + \epsilon_{idw} \tag{2.5}$$

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 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:

$$logP_{idw} = \alpha + \beta \hat{N}_{dw} + \gamma_w + \zeta_d + \epsilon_{idw} \quad (2.6)$$

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 2.3 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 2.1 presents heterogeneity in entry effects along different dimensions of market characteristics, which should be interpreted with caution. 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.

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 2.2 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. 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%).

2.6 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:[4]

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:

$$P_{dt} = c_{dt} - \theta \frac{\partial P_{dt}}{\partial Q_{dt}} \frac{Q_{dt}}{N_{dt}} \tag{2.7}$$

Assuming 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):[5]

$$E[P_T - P_C] = E\left[-\left(\frac{Q_T}{\sigma_T}\right)\left(\frac{\partial P_T}{\partial Q_T}\right) + \left(\frac{Q_C}{\sigma_C}\right)\left(\frac{\partial P_C}{\partial Q_C}\right)\right] \tag{2.8}$$

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

$$E[P_T - P_C] = E\left[\left(\frac{\delta(a-P_T)}{\sigma_T}\right) - \left(\frac{\delta(a-P_C)}{\sigma_C}\right)\right] \tag{2.9}$$

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 2.4 presents these simulations. The top panel presents the simulated effect of entry on one trader (the

[4]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.

[5]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.

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 2.9.[6]

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:

$$\overline{\Delta P^{EB}} = \frac{\sum_{s=1}^{10} \Delta P_S^{EB} N_s}{\sum_{s=1}^{10} N_S} \tag{2.10}$$

in which ΔP_S^{EB} is the change in price expected in a market of size S if entrants act according 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 2.3 presents a graphical version of this intuition. The bootstrapped distribution of the reduced form effect on log price is shown. 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. 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?

[6]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 2.9. 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 2.4. 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.

My analysis suggests we do. Take Equation 2.9 and solve for σ_T:

$$\sigma_T = \frac{\delta(a - E[P_T])}{E[P_T - P_C] + \frac{\delta(a - E[P_C])}{\sigma_C}} \tag{2.11}$$

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 2.11. The kernel density of the resulting σ_C and σ_T are displayed in Figure 2.4. 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 2.5 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.

2.7 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. 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 2.1) and in the final week of entry. The latter is suggestive of dynamic effects consistent with a breakdown in collusion in the final period (see Appendix A.5).

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.

2.8 Conclusion

In the context of developing county markets, encourage greater market entry may be a more feasible policy response to limited competition than regulation or state intervention. 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.

An ongoing project in Uganda tests the impact of one such technology. Joint with Craig McIntosh, this work evaluates the impact of a mobile marketplace that links farmers with buyers of agricultural commodities. The platform is designed to provide a national marketplace in which a larger number of potential trade partners can meet, thereby fostering greater competition and shortening output market chains. The randomized introduction of this technology is complemented by a massive price information campaign, in which we collect biweekly prices from 260 markets across Uganda and disseminate this information to farmers and intermediaries.

In addition to this main intervention, which is designed to reduce search costs and enhance competition, we also use several sub-experiments to assess the role of other potential barriers that may contribute to the shallowness of agricultural markets. One of these potential barriers is contractual risk. Traders must invest in transportation to rural villages without any guarantees that agreements regarding quantity or quality of available crops will be upheld upon arrival. Without legal recourse when agreements are breached, traders may find buying from remote areas too risky. They may instead choose to trade only with trusted farmers or brokers with whom they have had repeated interactions, and may simply forgo trade in remote areas entirely. We partner with one of Uganda?s largest private-sector brokerage companies to offer a guarantee to randomly selected buyers in the marketplace, in which they receive monetary compensation if they commit to a transaction and find the terms of the contract altered upon arrival. We test whether this reduction in risk encourages buyers to enter into agreements with new trading partners.

We also test the role that credit constraints play in hampering efficient trade. Again, with our brokerage company partner, we randomly offer loans to local intermediates and test whether the availability of credit allows them to deal in larger transactions than they would otherwise be unable to engage.

The cross-cutting design of this multi-pronged experiment is intended to disentangle the role of search costs, contractual risk, and credit constraints in limiting efficient arbitrage. Price surveys enable estimation of the impacts on price dispersion and market efficiency. Large-scale farmer and trader surveys shed light on the distribution of welfare effects from these interventions.

Tables and Figures

Figure 2.1: **H**eterogeneity 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.

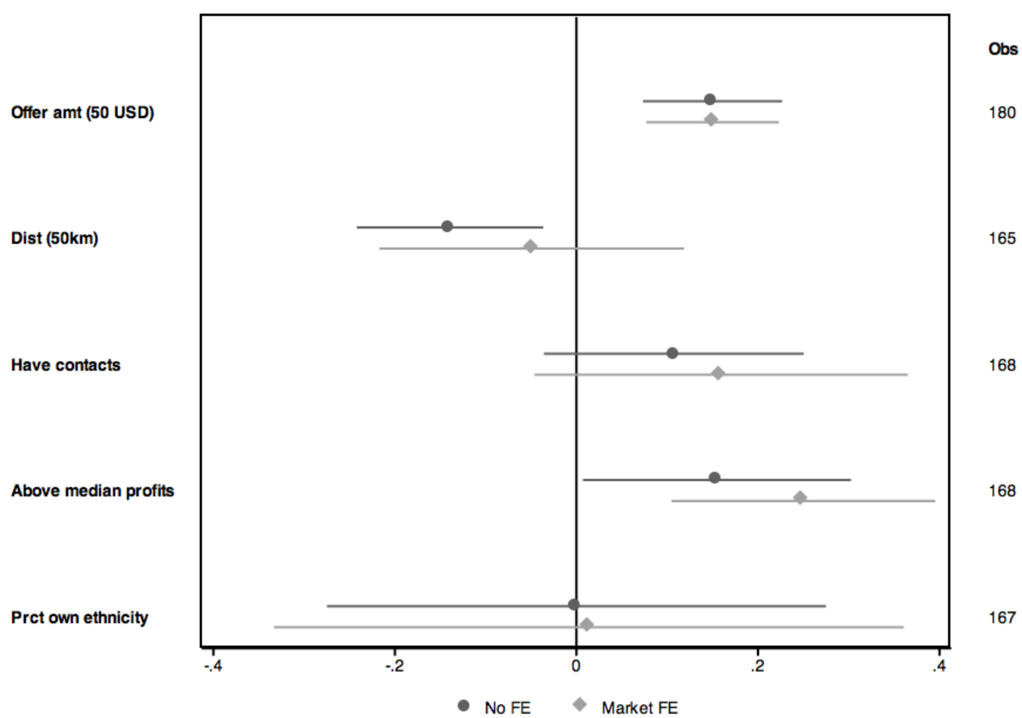

Figure 2.2: **H**eterogeneity in take-up and IV impact of entry by market characteristics. Number of traders divided into below and above median number (as counted in control periods). Tarmac is a dummy for whether the market is on a tarmac road or not. The "discuss (agree) price" variable is constructed 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 2.6, but with the sample restricted to the subgroup in question.

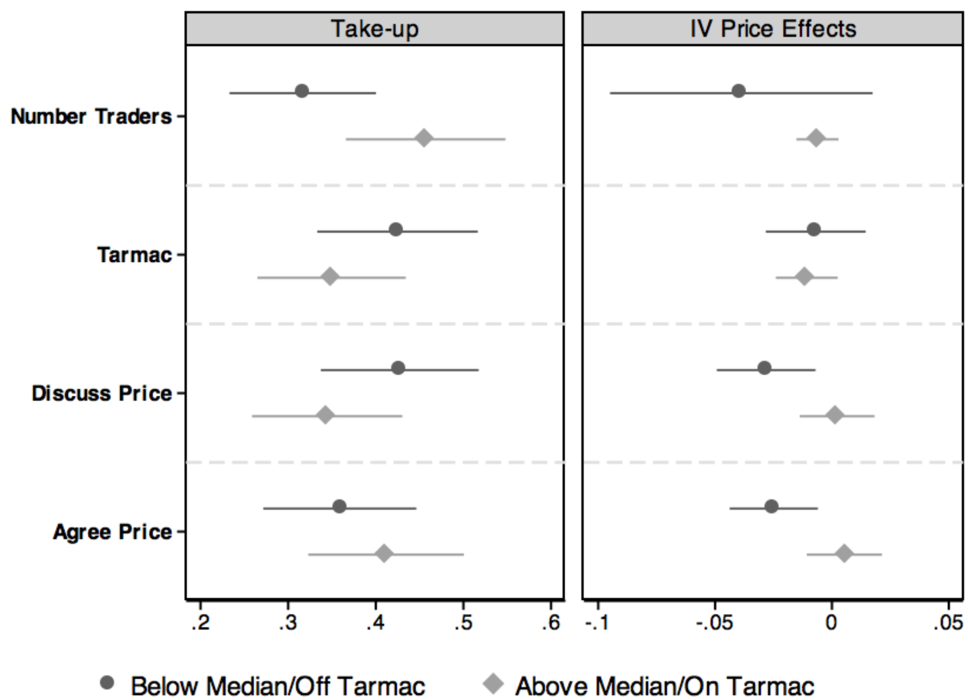

Figure 2.3: **P**redicted price change for various formed of behavior on the part of the entrant. Taking the point estimates on the demand estimation seriously, I plot the expected reduction in price one would have expected to see under different models of entrant behavior: entrant competing with the 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).

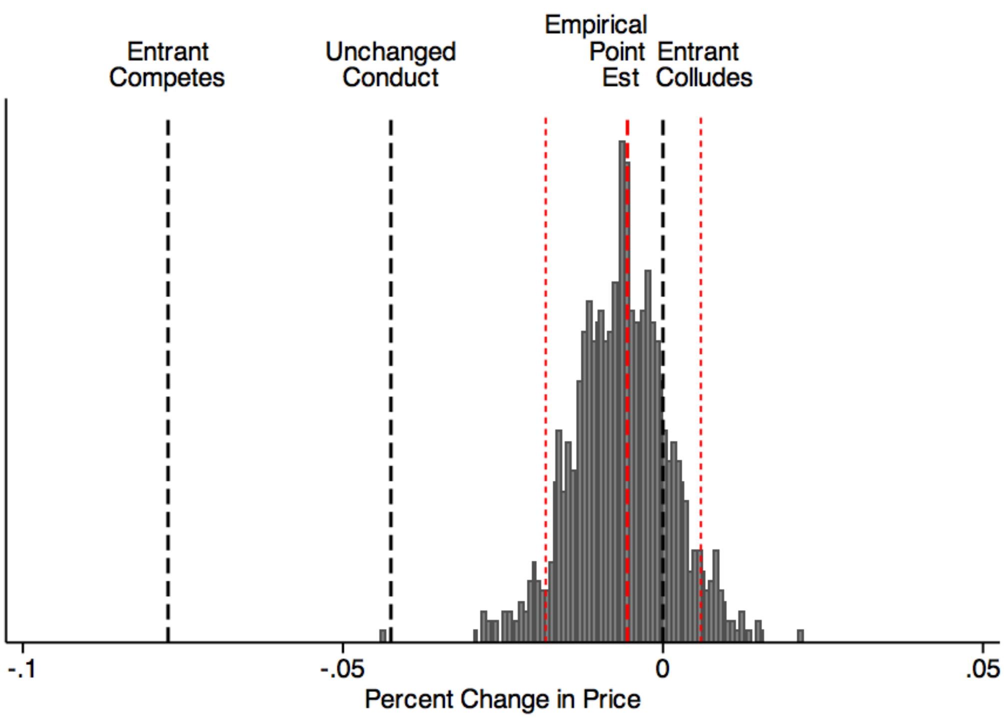

Figure 2.4: **K**ernel density of sigma control and sigma after entry. I take 1,000 bootstrapped samples of my entire data and, for each, calculate the baseline competitiveness parameter σ_C (as determined by the estimated pass-through rate int that sample and the estimate of demand curvature δ from 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).

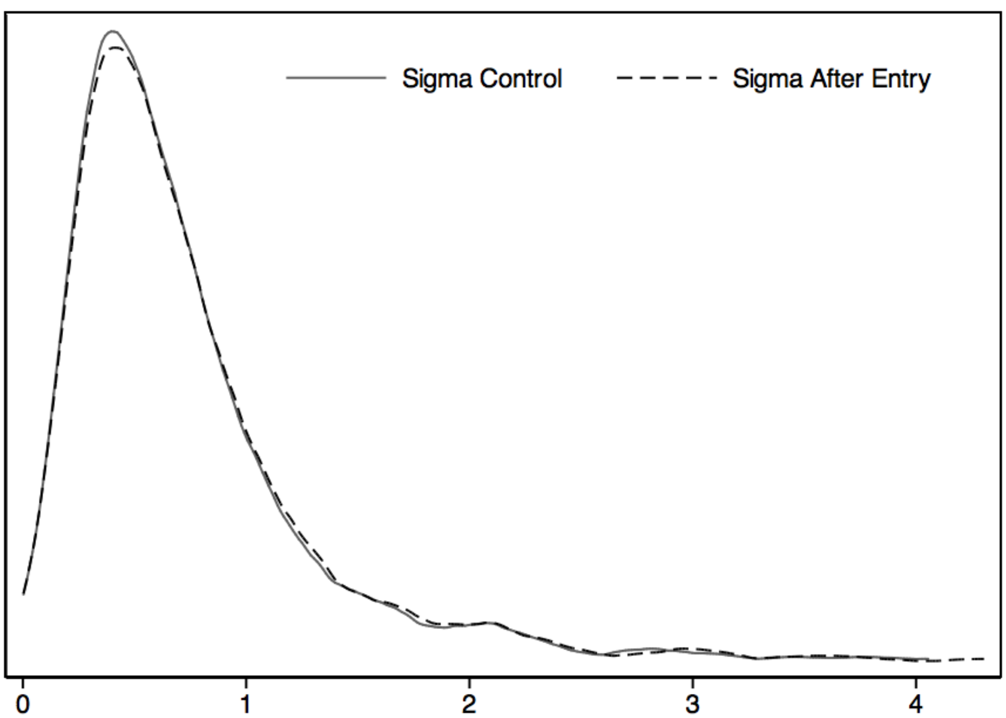

Figure 2.5: **D**istribution of change in sigma. The distribution of sigma T - sigma C from 1,000 bootstrapped samples is plotted. The mass is very close to 0 change (the outcome if the entrant colludes), with only slight increases in competition as seen by a slight shift to the right of the distribution. 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).

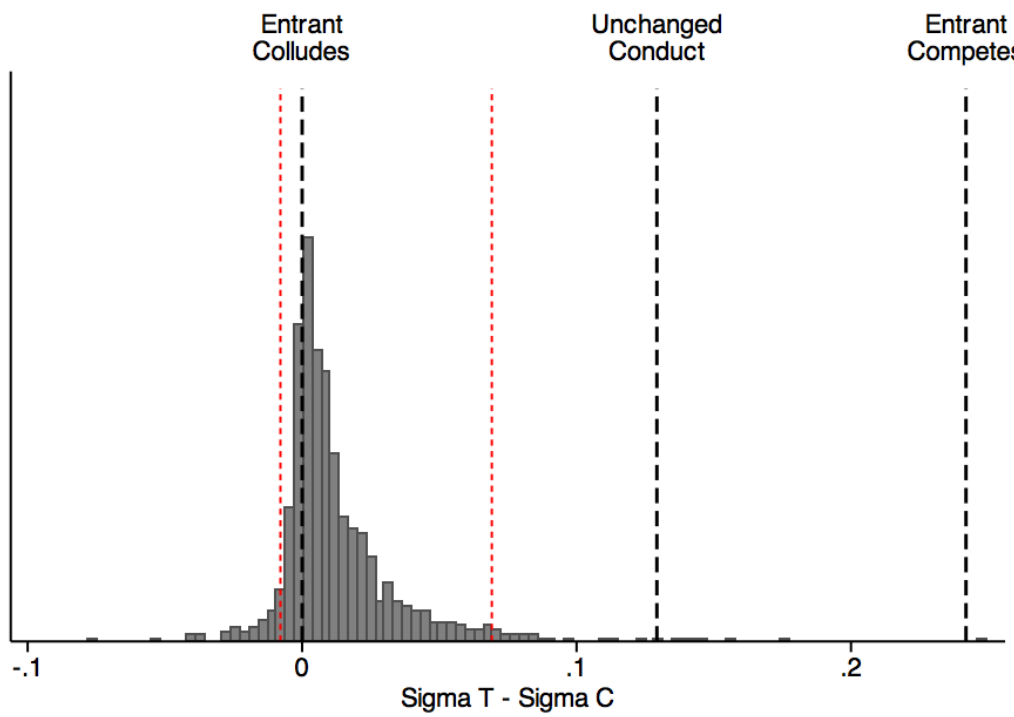

Table 2.1: **T**ake-up of Entry Offers. Offers ranged from 5,000-15,000 Kenyan shillings ($49-148 USD). Take-up rate = 1 if the trader *ever* took up an offer during any of the four weeks for which the offer was available. In a survey conducted during the offer phone call, traders reported the revenues 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 2.2: **T**ake-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 1.2), 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 2.3: **E**ffect 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 2.4: **S**imulated 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 **t**op panel presents simulated IV effects of impact of entry by one trader. The **m**iddle 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 **b**ottom panel simulated reduced form effects of the impact of entry using the average first-stage for all market-sizes.

Baseline Num Traders	**Num Mkts**	**Δ N**	**Conduct Unchanged** *Theta*	*Sigma*	*% Price Δ*	**Entrant Competes** *Theta*	*Sigma*	*% Price Δ*	**Entrant Colludes** *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

Chapter 3

Arbitrage: Seasonal Price Fluctuations, Storage, and the Returns to Credit in the Presence of General Equilibrium Effects

with Marshall Burke and Edward Miguel

3.1 Introduction

Large and regular seasonal price fluctuations are common in African agricultural markets. Grain prices in major markets typically rise by roughly 25-50% in the six months following harvest; increases as large as 100% are frequent in more isolated markets. Driving these fluctuations is puzzling behavior at the farmer level: despite having access to relatively cheap storage technologies, farmers tend to sell their crops immediately after harvesting, flooding the markets and driving prices down. In many areas, including our study area in Western Kenya, these same farmers later return to the market as consumers in the lean season, once prices have risen.

In this paper, we study the role that credit constraints play in farmers' inability to store grain and arbitrage these seasonal price fluctuations. Lacking access to credit or savings, farmers report selling their grain at low post-harvest prices to meet urgent cash needs (e.g., to pay school fees). To meet consumption needs later in the year, many then end up buying back grain from the market a few months after selling it, in effect using the maize market as a high-interest lender of last resort (Stephens and Barrett, 2011).

Working with a local agricultural microfinance NGO, we offer randomly selected smallholder maize farmers a loan at harvest, and study whether access to this loan improves their ability to use storage to arbitrage local price fluctuations, relative to a control group. We find that farmers offered the loan sell significantly less and purchase significantly more maize in the period immediately following harvest, and this pattern reverses during the period of high prices 6-9 months later. This change in the marketing behavior results in statistically significant and economically meaningful effects on farm revenue.

To test the robustness of these results, we replicate the experiment in two back-to-back years and find remarkably similar results on primary outcomes. We also run a long-run follow-up survey with respondents 1-2 years after credit was removed,[1], in order to test whether farmers are able to use the additional revenues earned from this loan product in one year to "save their way out" of credit constraints in future years. We see no long-run effects on inventories, sales, or revenues, lending little support to the idea that this particular short-run intervention can enable farmers to graduate out of poverty in this context.

Because storage-related changes in behavior could have effects on local prices in a setting of high regional transport costs, we vary the density of treated farmers across locations and track market prices at 50 local market points. Greater storage at the market level results in smoother prices over the season; in areas with high treatment density, prices immediately after harvest are significantly higher, while prices during the lean season are lower (although the latter not significantly so). Discernible price effects from such a localized shift in supply suggest that agricultural markets in the region are highly fragmented.

These general equilibrium effects alter the profitability of the loan. By dampening the arbitrage opportunity posed by season price fluctuations, greater treatment saturation results in diminished revenue impacts for treated individuals. We find that while treated farmers in high-density areas store significantly more than their control counterparts, they are not more profitable; the reduction in seasonal price dispersion in these area reduces the benefits of loan adoption. Conversely, treated farmers in low-density areas have both significantly higher inventories and significantly higher profits relative to control.

The result that general equilibrium effects quickly erode the profitability of arbitrage in highly fragmented markets has lessons for both policy and evaluation. In terms of policy, the quickly diminishing returns may be one reason (of many) why private sector lending institutions have not found it profitable to offer such arbitrage-oriented loans. At high saturation levels, banks may not be able to recoup a meaningful interest rate. Their optimal saturation rate (i.e. the number of loans they offer) may be quite low and therefore not justify the administrative costs of offering such a product. This does not mean, however, that the *socially* optimal saturation rate is low. It may still be socially optimal to offer widespread loans, despite the quickly diminishing incentives for the private-sector to do so. This is because a large portion of the benefits of loan expansion accrue to those who not receive the loans (in our experiment, the control group), a group from which private-sector institutions cannot capture returns but about whose utility a social planner cares.

The importance of general equilibrium effects in shaping the impact of these loans also has implications for evaluation. A simple individual-level randomized evaluation of such products – if conducted at a high saturation level – may lead us to conclude that the loan is

[1]While our NGO partner, One Acre Fund, is currently considering scaling up this loan product, in the period between the experiment and the long-run follow-up, no additional maize storage loans were offered.

ineffective, as treatment and control group may have statistically indistinguishable revenues. However, this is not due to a lack of impact of the loan, but rather to the fact that at high saturation levels the main gains from the arbitrage product come from the general equilibrium effects on prices, which accrue to both the treatment and control groups, rather than the private benefits to the treatment group alone. Take the most extreme case, in which saturation is high enough to produce perfect arbitrage, in which prices are equalized across seasons.[2] In this scenario, it will not matter when one sells; the treatment group, which sells later in the season, will be no better off than the control group, which sells earlier in the season. However, both benefit from smoother prices and overall higher revenues than they would in the absence of treatment, in which both groups sell early in the season for a lower price.[3]. When the benefits of an intervention come from general equilibrium effects that affect both treatment and control individuals, this intervention cannot be evaluated in partial equilibrium models with simple individual-level randomization.

This paper speaks to a rich literature on the importance of credit market imperfections in underdevelopment (Banerjee and Newman, 1993; Galor and Zeira, 1993; Banerjee and Duflo, 2010). Findings in this rapidly growing literature have been remarkably heterogenous. Studies that provide cash grants to households and to existing small firms suggest high rates of return to capital in some settings but not in others.[4] Further, experimental evaluations of traditional microcredit products (small loans to poor households) have generally found that individuals randomly provided access to these products are subsequently no more productive on average that those not given access, but that subsets of recipients often appear to benefit.[5]

Why do we find positive effects on firm profitability when many other experimental studies on microcredit do not? These studies have offered a number of explanations as to why improved access to capital does not appear beneficial on average. First, many small businesses or potential micro-entrepreneurs simply might not actually face profitable investment opportunities (Banerjee et al., 2013; Fafchamps et al., 2013; Karlan, Knight and Udry, 2012; Banerjee, 2013).[6] Second, profitable investment opportunities could exist but established or potential microentrepreneurs might lack either the skills or ability to channel capital towards

[2]For the sake of simplicity, we ignore the role of interest rates here.

[3]On the other hand, individual-level randomization at a low saturation may *overstate* the benefits to the treatment group compared to a program rolled out at scale (and may still miss any ? albeit smaller – benefits that accrue to the control group). The lesson remains; in situations in which the level and distribution of benefits are shaped by general equilibrium effects, proper measurement of these features requires a study design that randomizes the intensity of these general equilibrium effects (e.g. random saturation)

[4]Studies finding high returns to cash grants include De Mel, McKenzie and Woodruff (2008); McKenzie and Woodruff (2008); Fafchamps et al. (2013); Blattman, Fiala and Martinez (2013). Studies finding much more limited returns include Berge, Bjorvatn and Tungodden (2011) and Karlan, Knight and Udry (2012).

[5]Experimental evaluations of microcredit include Attanasio et al. (2011); Crepon et al. (2011); Karlan and Zinman (2011); Banerjee et al. (2013); Angelucci, Karlan and Zinman (2013). See Banerjee (2013) and Karlan and Morduch (2009) for nice recent reviews of these literatures.

[6]For example, many microenterprises might have low efficient scale and thus little immediate use for additional investment capital, with microentrepreneurs then preferring to channel credit toward consumption instead of investment. Relatedly, marginal returns to investment might be high but total returns low, with the entrepreneur making the similar decision that additional investment is just not worth it.

these investments - e.g. if they lack managerial skills (Berge, Bjorvatn and Tungodden, 2011; Bruhn, Karlan and Schoar, 2012), or if they face problems of self-control or external pressure that redirect cash away from investment opportunities (Fafchamps et al., 2013). Third, typical microcredit loan terms require that repayment begin immediately, and this could limit investment in illiquid but high-return business opportunities (Field et al., 2012). Finally, as described above, general equilibrium effects of credit expansion could alter individual-level treatment effect estimates in a number of ways, potentially shaping outcomes for treated individuals. [7] This is a recognized but unresolved problem in the experimental literature on credit, and few experimental studies have been explicitly designed to quantify these effects.[8]

All of these factors likely help explain why our results diverge from existing estimates. Unlike most of the settings examined in the literature, using credit to "free up" storage for price arbitrage does not require starting or growing a business among this population of farmers, is neutral to the scale of farm output, does not appear to depend on entrepreneurial skill (all farmer have stored before, and all are very familiar with local price movements), and does not require investment in a particularly illiquid asset (inventories are kept in the house and can be easily sold). Farmers do not even have to sell grain to benefit from credit in this context: a net-purchasing farm household facing similar seasonal cash constraints could use credit and storage to move purchases from times of high prices to times of low prices.

Furthermore, our results also suggest that – at least in our rural setting – treatment density matters and market-level spillovers can substantially shape individual-level treatment effect estimates. Whether these GE also influenced estimated treatment effects in more urban settings is unknown, although there is some evidence that spillovers do matter for microenterprises who directly compete for a limited supply of inputs to production.[9] In any case, our results suggest that explicit attention to GE effects in future evaluations of credit market interventions is likely warranted.

Beyond contributing to the experimental literature on microcredit, our paper is closest to a number of recent papers that examine the role of borrowing constraints in households' storage decisions and seasonal consumption patterns. Using secondary data from Kenya, Stephens and Barrett (2011) also suggest that credit constraints substantially alter small-

[7] While this issue may be particularly salient in our context of a loan explicitly designed to enable arbitrage, it is by no means unique to our setting. Any enterprise operating in a small, localized market or in a concentrated industry may face price responses to shifts in own supply. Credit-induced expansion may therefore be less profitable than it would be in more integrated market or less concentrated industries.

[8] For instance, Karlan, Knight and Udry (2012) conclude by stating, "Few if any studies have satisfactorily tackled the impact of improving one set of firms' performance on general equilibrium outcomes.... This is a gaping hole in the entrepreneurship development literature." Indeed, positive spillovers could explain some of the difference between the experimental findings on credit, which suggest limited effects, and the estimates from larger-scale natural experiments, which tend to find positive effects of credit expansion on productivity – e.g. Kaboski and Townsend (2012). Acemoglu (2010) uses the literature on credit market imperfections to highlight the understudied potential role of GE effects in broad questions of interest to development economists.

[9] See De Mel, McKenzie and Woodruff (2008) and their discussion of returns to capital for firms in the bamboo sector, all of whom in their setting compete over a limited supply of bamboo.

holder farmers' marketing and storage decisions, and Basu and Wong (2012) show that allowing farmers to borrow against future harvests can substantially increase lean-season consumption. As in these papers, our results show that when borrowing and saving are difficult, households turn to increasingly costly ways to move consumption around in time. In our particular setting, credit constraints combined with post-harvest cash needs cause farmers to store less than they would in an unconstrained world, lowering farm profits even in a year when prices don't rise much. In this setting, even a relatively modest expansion of credit affects local market prices, to the apparent benefit of those with and without access to this credit.

Finally, our results speak to an earlier literature showing how credit market imperfections can combine with other features of economies to generate observed broad-scale economic patterns (Banerjee and Newman, 1993; Galor and Zeira, 1993). These earlier papers showed how missing markets for credit, coupled with an unequal underlying wealth distribution, could generate large-scale patterns of occupational choice. we show that missing markets for credit combined with climate-induced seasonality in rural income can help generate widely-observed seasonal price patterns in rural grain markets, patterns that appear to further worsen poor households' abilities to smooth consumption across seasons. That expansion of credit access appears to help reduce this price dispersion suggests an under-appreciated but likely substantial additional benefit of credit expansion in rural areas.

The remainder of the paper proceeds as follows. Section 3.2 describes the setting and the experiment. Section 3.3 describes our data, estimation strategy, and pre-analysis plan. Section 3.4 presents baseline estimates ignoring the role of general equilibrium effects. Section 3.5 presents the market level effects of the intervention, and shows how these affect individual-level estimates. Section 3.6 concludes.

3.2 Setting and experimental design

Arbitrage opportunities in rural grain markets

Seasonal fluctuations in prices for staple grains appear to offer substantial intertemporal arbitrage opportunities, both in our study region of East Africa as well as in other parts of Africa and elsewhere in the developing world. While long term price data unfortunately do not exist for the small markets in very rural areas where our experiment takes place, price series are available for major markets throughout the region. Average seasonal price fluctuations for maize in available markets are shown in Figure 3.1. Increases in maize prices in the six to eight months following harvest average roughly 25-50% in these markets, and these increases appear to be a lower bound on seasonal price increases reported elsewhere in Africa.[10]

[10]For instance, Barrett (2008) reports seasonal rice price variation in Madagascar of 80%, World Bank (2006) reports seasonal maize price variation of about 70% in rural Malawi, and Aker (2012) reports seasonal variation in millet prices in Niger of 40%.

These increases also appear to be a lower bound on typical increase observed in the smaller markets in our study area, which (relative to these much larger markets) are characterized with much smaller "catchments" and less outside trade. We asked farmers at baseline to estimate average monthly prices for either sales or purchases of maize at their local market point over the last five years, and as shown in the left panel of Figure 3.2, they reported a typical doubling in price between September (the main harvest month) and the following June. In case farmers were somehow mistaken or overoptimistic, we asked the same question of the local maize traders that can typically be found in these market points. These traders report very similar average price increases: the average reported increase between October and June across traders was 87% (with a 25th percentile of 60% increase and 75th percentile of 118% - results available on request).

Farmers do not appear to be taking advantage of these apparent arbitrage opportunities. Figure C.1 shows data from two earlier pilot studies conducted either by One Acre Fund (in 2010/11, with 225 farmers) or in conjunction with One Acre Fund (in 2011/12, with a different sample of 700farmers). These studies tracked maize inventories, purchases, and sales for farmers in our study region. In both years, the median farmer exhausted her inventories about 5 months after harvest, and at that point switched from being a net seller of maize to a net purchaser as shown in the right panels of the figure. This was despite the fact that farmer-reported sales prices rose by more than 80% in both of these years in the nine months following harvest.

Why are farmers not using storage to sell at higher prices and purchase at lower prices? Our experiment is primarily be designed to test the role of credit constraints in shaping storage and marketing decisions. In extensive focus groups with farmers prior to our experiment, credit constraints were the (unprompted) explanation given by the vast majority of these farmers as to why they were not storing and selling maize at higher prices. In particular, because nearly all of these farm households have school aged kids, and a large percentage of a child's school fees are typically due in the few months after harvest (prior to January enrollment), many farmers report selling much of their harvest to pay these fees. Indeed, many schools in the area will accept in-kind payment in maize during this period. Farmers also report having to pay other bills they have accumulated throughout the year during the post-harvest period.

Further, as with poor households throughout much of the world, these farmers appear to have very limited access to formal credit. Only eight percent of households in our sample reported having taking a loan from a bank in the year prior to the baseline survey. Informal credit markets also appear relatively thin, with less than 25% of farmers reporting having given or received a loan from a moneylender, family member, or friend in the 3 months before the baseline.

Absent other means of borrowing, and given these various sources of "non-discretionary" consumption they report facing in the post-harvest period, farmers end up liquidating rather than storing. Furthermore, a significant percentage of these households end up buying back maize from the market later in the season to meet consumption needs, and this pattern of "selling low and buying high" directly suggests a liquidity story: farmers are in effect taking a high-interest quasi-loan from the maize market (Stephens and Barrett, 2011). Baseline data indicate that 35% of our sample both bought and sold maize during the previous crop year (September 2011 to August 2012), and that over half of these sales occurred before January (when prices were low). 40% of our sample reported only purchasing maize over this period, and the median farmer in this group made all of their purchases after January. Stephens and Barrett (2011) report very similar patterns for other households in Western Kenya during an earlier period.

Nevertheless, there could be other reasons beyond credit constraints why farmer are not taking advantage of apparent arbitrage opportunities. The simplest explanations are that farmers do not know about the price increases, or that it's actually not profitable to store – i.e. arbitrage opportunities are actually much smaller than they appear because storage is costly. These costs could come in the form of losses to pests or moisture-related rotting, or they could come in the form of "network losses" to friends and family, since maize is stored in the home and is visible to friends and family, and there is often community pressure to share a surplus. Third, farmers could be highly impatient and thus unwilling to move consumption to future periods in any scenario. Finally, farmers might view storage as too risky an investment.

Evidence from pilot and baseline data, and from elsewhere in the literature, argues against a few of these possibilities. We can immediately rule out an information story: as shown in Figure 3.2 and discussed above, all farmers know exactly what prices are doing, and all expect prices to rise substantially throughout the year.[11] Second, pest-related losses appear surprisingly low in our setting, with farmers reporting losses from pests and moisture-related rotting of less than 5% for maize stored for six to nine months. Similarly, the fixed costs associated with storing for these farmers are small and have already been paid: all farmers store at least some grain (note the positive initial inventories in Figure C.1), and grain in simply stored in the household or in small sheds previously built for the purpose. Third, existing literature shows that for households that are both consumers and producers of grain, aversion to price risk should motivate *more* storage rather than less: the worst state of the world for these households is a huge price spike during the lean season, which should motivate "precautionary" storage (Saha and Stroud, 1994; Park, 2006). Fourth, while we cannot rule out impatience as a driver of low storage rates, extremely high discount rates would be needed to rationalize this behavior in light of the expected nine-month doubling of prices. Furthermore, farm households are observed to make many other investments with payouts far in the future (e.g. school fees), meaning that rates of time preference would also have to differ substantially across investments and goods.

Costs associated with network-related losses appear a more likely explanation for an unwillingness to store substantial amounts of grain. Existing literature suggests that com-

[11]The mean across farmers for all three reported prices (the historical purchase price, the historical sales price, and the expected sales price) is a 115-134% increase in prices. For the expected sales price over the ensuing nine months after the September 2012 baseline, the 5th, 10th, and 25th percentiles of the distribution are a 33%, 56%, and 85% increase, respectively, suggesting that nearly all farmers in our sample expect substantial price increases.

munity pressure is one explanation for limited informal savings (Dupas and Robinson, 2013; Brune et al., 2011), and in focus groups farmers often told us something similar about stored grain (itself a form of savings). As described below, our main credit intervention might also provide farmers a way to shield stored maize from their network. To further test this hypothesis, we add an additional treatment arm to determine whether this shielding effect is substantial on its own.

Experimental design

Our study sample is drawn from existing groups of One Acre Fund (OAF) farmers in Webuye and Matete districts in Western Kenya. OAF is a microfinance NGO that makes in-kind, joint-liability loans of fertilizer and seed to groups of farmers, as well as providing training on improved farming techniques. OAF group sizes typically range from 8-12 farmers, and farmer groups are organized into "sublocations" – effectively clusters of villages that can be served by one OAF field officer. OAF typically serves 20-30% of farmers in a given sublocation.

The Year 1 sample consists of 240 existing OAF farmer groups drawn from 17 different sublocations in Webuye district, and our total sample size at baseline was 1589 farmers. The Year 2 sample follows the same OAF groups as Year 1; however there was substantial shifting of the individual members that form these groups, and therefore some Year 1 farmers drop out of our Year 2 sample, and other farmers are new to our Year 2 sample.[12] Ultimately, of the 1,019 individuals in our Year 2 sample, 602 are drawn from the Year 1 sample and 417 are new to the sample.

There are two main levels of randomization. First, we randomly divided the 17 sublocations in our sample into 9 "high" treatment intensity sites and 8 "low" treatment density sites. We fixed the "high" treatment density at 80% (meaning 80% of groups in the sublocation would be offered a loan), and then determined the number of groups that would be needed in the "low" treatment sites in order to get our total number of groups to 240 (what the power calculations suggested we needed to be able to discern meaningful impacts at the individual level). This resulted in a treatment intensity of 40% in the "low" treatment-intensity sites, yielding 171 total treated groups in the high intensity areas and 69 treated groups in the low intensity areas in Year 1. Then, within each sublocation, groups were randomized into treatment or control, at a rate determined by the saturation rate of their sublocation.

In Year 1, the group-level randomization was stratified at the sublocation level, and then further stratified based on whether group-average OAF loan size in the previous year was above or below the sample median (data from the previous year were available from

[12] Shifting of group members is a function of several factors, including whether farmers wished to participate in the overall OAF program from year to year. There was some (small) selective attrition based on treatment status in Year 1; treated individuals were 10 percentage points more likely to return to the Year 2 sample than control individuals (significant at 1%). This does slightly alter the composition of the Year 2 sample (see Table C.30 and Section C.8), but because Year 2 treatment status is stratified by Year 1 treatment status (as will be described below), it does not alter the internal validity of the Year 2 results.

administrative data). In Year 2, we maintained the same saturation treatment status at the sublocation level,[13] but re-randomized groups into treatment and control, stratifying on their treatment status from Year 1.[14]

In Year 1, there was a third level of randomization pertaining to the timing of the loan offer. In focus groups run prior to the experiment, farmers were split on when credit access would be most useful, with some preferring cash immediately at harvest, and others preferring it a few months later timed to coincide with when school fees were due (the latter preferences suggesting that farmers may be sophisticated about potential difficulties in holding on to cash between the time it was disbursed and the time it needed to be spent). In order to test the importance of loan timing, in Year 1, a random half of the group received the loan in October (immediately following harvest), while the other half received the loan in January (immediately before school fees are due). As will be described in Section 3.4, results from Year 1 suggested that the earlier loan was more effective, and therefore in Year 2 we only offer the earlier timed loan to the full sample (though due to administrative delays, the actual loan was disbursed in November in Year 2).

Although all farmers in each loan treatment group were offered the loan, we follow only a randomly selected 6 farmers in each loan group, and a randomly selected 8 farmers in each of the control groups.

Loan offers were announced in September in both years. To qualify for the loan, farmers had to commit maize as collateral, and the size of the loan they could qualify for was a linear function of the amount they were willing to collateralize (capped at 7 bags in Year 1 and 5 bags in Year 2). In Year 1, to account for the expected price increase, October bags were valued at 1500Ksh, and January bags at 2000Ksh. In Year 2, bags were valued at 2500Ksh. Each loan carried with it a "flat" interest rate of 10%, with full repayment due after nine months.[15,16] These loans were an add-on to the existing in-kind loans that OAF clients received, and OAF allows flexible repayment of both – farmers are not required to repay anything immediately.

Collateralized bags of maize were tagged with a simple laminated tag and zip tie. When we mentioned in focus groups the possibility of OAF running a harvest loan program, and described the details about the collateral and bag tagging, many farmers (unprompted) said that the tags alone would prove useful in shielding their maize from network pressure: "branding" the maize as committed to OAF, a well-known lender in the region, would allow

[13]Such that, for example, if a sublocation was a high intensity sublocation in Year 1 it remained a high intensity sublocation in Year 2.

[14]This was intended to result in randomized duration of treatment – either zero years of the loan, one year of the loan, or two years – however, due to selective attrition of the Year 1 sample based on treatment status, duration of loan treatment is no longer entirely random.

[15]Annualized, this interest rate is slightly lower than the 16-18% APR charged on loans at Equity Bank, the main rural lender in Kenya.

[16]For example, a farmer who committed 5 bags when offered the October loan in Year 1 would receive 5*1500 = 7500Ksh in cash in October (∼$90 at current exchange rates), and would be required to repay 8250Ksh by the end of July.

them to credibly claim that it could not be given out.[17] Because tags could represent a meaningful treatment in their own right, we wished to separate the effect of the credit from any effect of the tag, and therefore in the Year 1 study offered a separate treatment arm in which groups received only the tags.[18]

Finally, because self- or other-control problems might make it particularly difficult to channel cash toward productive investments in settings where there is a substantial time lag between when the cash is delivered and when the desired investment is made, in Year 1, we also cross-randomized a simple savings technology that had shown promise in a nearby setting (Dupas and Robinson, 2013). In particular, a subset of farmers in each loan treatment group in Year 1 were offered a savings lockbox (a simple metal box with a sturdy lock) which they could use as they pleased. While such a savings device could have other effects on household decision making, our hypothesis was that it would be particularly helpful for loan clients who received the cash before it was needed.

The tags and lockbox treatments were randomized at the individual level during Year 1 (these treatments were dropping Year 2). Using the sample of individuals randomly selected to be followed in each group, we stratified individual level treatments by group treatment assignment and by gender. So, for instance, of all of the women who were offered the October Loan and who were randomly selected to be surveyed, one third of them were randomly offered the lockbox (and similarly for the men and for the January loan). In the control groups, in which we were following 8 farmers, 25% of the men and 25% of the women were randomly offered the lockbox, with another 25% each being randomly offered the tags (Ct). The study design allows identification of the individual and combined effects of the different treatments, and our approach for estimating these effects is described below.

3.3 Data and estimation

In August/September 2012 (prior to the Year 1 experiment), a baseline was conducted with the entire Year 1 sample. The baseline survey collected data on farming practices, on storage costs, on maize storage and marketing over the previous crop year, on price expectations for the coming year, on food and non-food consumption expenditure, on household borrowing, lending, and saving behavior, on household transfers with other family members and neighbors, on sources of non-farm income, on time and risk preferences, and on digit span recall.

We then undertook three follow-up rounds over the ensuing 12 months, spanning the spring 2013 "long rains" planting (the primary growing season) and concluding just prior to the 2013 long rains harvest (which occurs August-September). The multiple follow-up

[17]Such behavior is consistent with evidence from elsewhere in Africa that individuals take out loans or use commitment savings accounts mainly as a way to demonstrate that they have little to share (Baland, Guirkinger and Mali, 2011; Brune et al., 2011).

[18]This is of course not perfect – there could be an interaction between the tag and the loan – but we did not have the sample size to do the full 2 x 2 design to isolate any interaction effect.

rounds were motivated by three factors. First, a simple inter-temporal model of storage and consumption decisions suggests that while the loan should increase total consumption across all periods, the per-period effects could be ambiguous – meaning that consumption throughout the follow-up period needs to be measured to get at overall effects. Second, because nearly all farmers deplete their inventories before the next harvest, inventories measured at a single follow-up one year after treatment would likely provide very little information on how the loan affected storage and marketing behavior. Finally, as shown in McKenzie (2012), multiple follow-up measurements on noisy outcomes variables (e.g consumption) has the added advantage of increasing power. A similar schedule of three follow-up rounds over 12 months were run in Year 2.[19] The follow-up surveys tracked data on storage inventory, maize marketing behavior, consumption, and other credit and savings behavior. Follow-up surveys also collected information on time preferences and on self-reported happiness.

In order to explore the long-run effects of the loan, we also ran a Long-Run Follow-Up (LRFU) survey from November-December 2015. This was two (one) years following loan repayment for the Year 1 (Year 2) treatment group. This survey followed up on the entire Year 2 sample (1,091 individuals) and a random subset of the Year 1 only sample (another 481 individuals), for a total sample of 1500 individuals. The survey collected information on maize harvests, sales, purchases, and revenues from 2014-2015 (broken down by harvest and lean season). It also collected data on farm inputs (labor and capital), food consumption and expenditure, household consumption, educational expenditure and attendance among children, non-farm employment and revenues, and a self-reported happiness measure. We were able to track 91.5% of the intended sample. There is no differential attrition based on Year 2 treatment status. While there is some suggestive evidence of differential attrition based on Year 1 treatment status (being treated in Year 1 is associated with 3 percentage point increase in the likelihood of being found in the long-run follow up survey, significant at 10%), this is partially driven by the fact that Year 1 treated individuals were more likely to be in the Year 2 sample (and therefore had been more recently in touch with our survey team). After controlling for whether an individual was present in the Year 2 sample, Year 1 treatment status is no longer significantly correlated with attrition.

In addition to farmer-level surveys, we also collected monthly price surveys at 52 market points in the study area. The markets were identified prior to treatment based on information from local OAF staff about the market points in which client farmers typically buy and sell maize. Data collection for these surveys began in November 2012 and continued through December 2015.

Finally, loan repayment data from OAF administrative records that was generously

[19] Because the Year 2 experiment was meant to follow the sample sample as Year 1, a second baseline was not run prior to Year 2. However, as described in Section 3.2, due to administrative shifts in farmer group composition, 417 of the 1,019 individuals in the Year 2 sample were new to the study. For these individuals, we do not have baseline data (there was insufficient time between receiving the updated administrative records for Year 2 groups and the disbursal of the loan to allow for a second baseline to be run). Therefore, balance tables can only be run with the sample that was present in Year 1. Because the loan offer was randomized, however, this should not affect inference regarding the impacts of the loan.

shared by OAF.

Table 3.1 shows summary statistics for a range of variables at baseline, and shows balance of these variables across the three main loan treatment groups. Groups are well balanced, as would be expected from randomization. Table C.27 shows the analogous table comparing individuals in the high- and low-treatment-density areas; samples appear balanced on observables here as well. Attrition was also relatively low across our survey rounds. In Year 1, overall attrition was 8%, and not significantly different across treatment groups (8% in the treatment group and 7% in the control). In Year 2, overall attrition was 2% (in both treatment and control, with no significant difference). There was some (small) selective attrition the Year 1 to the Year 2 sample based on Year 1 treatment status; treated individuals were 10 percentage points more likely to return to the Year 2 sample than control individuals (significant at 1%). This does slightly alter the composition of the Year 2 sample (see Table C.30), but because Year 2 treatment status is stratified by Year 1 treatment status, it does not alter the internal validity of the Year 2 results. Appendix C.8 explores this further.

Pre-analysis plan

To limit both risks and perceptions of data mining and specification search (Casey, Glennerster and Miguel, 2012), we specified and registered a pre-analysis plan (PAP) for Year 1 prior to the analysis of any follow-up data.[20] The Year 2 analysis follows a near identical analysis plan. Both the PAP and the complete set of results are available upon request.

We deviate significantly from the PAP in one instance: as described below, it became clear that our method for estimating market-level treatment effects specified in the pre-analysis plan could generate biased estimates, and here we pursue an alternate strategy that more directly relies on the randomization. In two other instances we add to the PAP. First, in addition to the regression results specified in the PAP, we also present graphical results for many of the outcomes. These results are just based on non-parametric estimates of the parametric regressions specified in the PAP, and are included because they clearly summarize how treatment effects evolve over time, but since they were not mentioned in the PAP we mention them here. Second, we failed to include in the PAP the (obvious) regressions in which the individual-level treatment effect is allowed to vary by the sublocation-level treatment intensity.

Estimation of treatment effects

In all analyses, we present results separately by year and pooled across years. Because the Year 2 replication produced results that are quantitatively quite similar to the Year 1 results for most outcomes, we rely on the pooled results as our specification of primary interest. However, for the sake of transparency and – for the outcomes in which the two years' results diverse – comparison, we report both.

[20]The pre-analysis plan is registered here: https://www.socialscienceregistry.org/trials/67, and was registered on September 6th 2013.

We have three main outcomes of interest: inventories, maize net revenues, and consumption. Inventories are the number of bags the household had in their maize store at the time of the each survey. This amount is visually verified by our enumeration team, and so is likely to be measured with very little error. We define maize net revenues as the value of all maize sales minus the value of all maize purchases, and minus any additional interest payments made on the loan for individuals in the treatment group. We call this "net revenues" rather than "profits" since we likely do not observe all costs; nevertheless, costs are likely to be very similar across treatment groups (fixed costs were already paid, and variable costs of storage are very low). The values of sales and purchases were based on recall data over the period between each survey round. Finally, we define consumption as the log of total per capita household expenditure over the 30 days prior to each survey. For each of these variables we trim the top and bottom 0.5% of observations, as specified in the pre-analysis plan.

Letting T_{jy} be an an indicator for whether group j was assigned to treatment in year y, and y_{ijry} as the outcome of interest for individual i in group j in round $r \in (1, 2, 3)$ in year y. Our main specification pools data across follow-up rounds 1-3 (and for the pooled specification, across years):

$$Y_{ijry} = \alpha + \beta T_{jy} + X_{ijry}\eta_{ry} + \varepsilon_{ijry} \tag{3.1}$$

The coefficient β estimates the Intent-to-Treat and, with round-year fixed effects η_{ry}, is identified from within-round variation between treatment and control groups. β can be interpreted as the average effect of being offered the loan product across follow-up rounds. Standard errors are clustered at the group level.

To absorb additional variation in the outcomes of interest, we also control for survey date in the regressions. Each follow-up round spanned 3+ months, meaning that there could be (for instance) substantial within-round drawdown of inventories. Inclusion of all of this control should help to make our estimates more precise without changing point estimates.

The assumption in (3.1) is that treatment effects are constant across rounds. In our setting, there are reasons why this might not be the case. In particular, if treatment encourages storage, one might expect maize revenues to be *lower* for the treated group immediately following harvest, as they hold off selling, and *greater* later on during the lean season, when they release their stored grain. To explore whether treatment effects are constant across rounds, we estimate:

$$Y_{ijry} = \sum_{r=1}^{3} \beta_r T_{jy} + X_{ijry} + \eta_{ry} + \varepsilon_{ijry} \tag{3.2}$$

and test whether the β_r are the same across rounds (as estimated by interacting the treatment indictor with round dummies). Unless otherwise indicated, we estimate both (3.1) and (3.2) for each of the hypotheses below.

To quantify market level effects of the loan intervention, we tracked market prices at 52 market points throughout our study region, and we assign these markets to the nearest sublocation. We begin by estimating the following linear model[21]:

$$y_{msty} = \alpha + \beta_1 H_s + \beta_2 month_t + \beta_3 (H_s * month_t) + \varepsilon_{mst} \tag{3.3}$$

where y_{mst} represents the maize sales price at market m in sublocation s in month t in year y. H_s is a dummy for if sublocation s is a high-intensity sublocation, and $month_t$ is a time trend (in each year, Nov = 1, Dec = 2, etc). If access to the storage loan allowed farmers to shift purchases to earlier in the season or sales to later in the season, and if this shift in marketing behavior was enough to alter supply and demand in local markets, then our prediction is that $\beta_1 > 0$ and $\beta_3 < 0$, i.e. that prices in areas with more treated farmers are higher after harvest but lower later in the year.

While H_s is randomly assigned, and thus the number of treated farmers in each sublocation should be orthogonal to other location-specific characteristics that might also affect prices (e.g. the size of each market's catchment), we are only randomizing across 17 sublocations. This relatively small number of clusters could present problems for inference (Cameron, Gelbach and Miller, 2008). We begin by clustering errors at the sublocation level when estimating (3.3). Future versions of the will also report standard errors estimated using both the wild bootstrap technique described in Cameron, Gelbach and Miller (2008), and the randomization inference technique (e.g. as used by Cohen and Dupas (2010)).

To understand how treatment density affects individual-level treatment effects, we estimate Equations 3.1 and 3.2, interacting the individual-level treatment indicator with the treatment density dummy. The pooled equation is thus:

$$Y_{ijsry} = \alpha + \beta_1 T_{jy} + \beta_2 H_s + \beta_3 (T_{jy} * H_s) + X_{ijry} + \eta_{ry} + \varepsilon_{ijsry} \tag{3.4}$$

If the intervention produces enough individual level behavior to have market effects, we predict that $\beta_3 < 0$ and perhaps that $\beta_2 > 0$ - i.e. treated individual in high-density areas do worse than in low density areas, and control individuals in high density areas do better (due to higher initial prices at which they'll be selling their output). As in Equation 3.3, we will report results with errors clustered at the sublocation level.

For long-run effects, we first estimate the following regression for each year separately:

$$Y_{ij} = \alpha + \beta T_{jy} + \varepsilon_{ij} \tag{3.5}$$

in which Y_{ij} is the outcome of interest for individual i in group j. The sample is restricted to those who were in the Year y study.

We further also run the following specification:

$$Y_{ij} = \alpha + \beta_1 T_{j1} + \beta_2 T_{j2} + \beta_3 T_{j1} * T_{j2} + \varepsilon_{ij} \tag{3.6}$$

[21] This estimating equation is slightly different than what was proposed in the pre-analysis plan. As was energetically pointed out to the authors during a seminar presentation at Berkeley after the pre-analysis plan had been registered, the proposed estimating equation for quantifying market level effects (which relied on counting up the number of treated farmers) could produce biased estimates because we are in practice unable to control for the total number of farmers in the area. Using the randomization dummy avoids this worry.

in which T_{j1} is a dummy for being an in treated group in year 1, T_{j12} is a summary for being in a treated group in year 2, and $T_{j1} * T_{j2}$ is an interaction term for being in a group that was treated in both years. The sample is restricted to those who were in the study for both years. Because of this sample restriction, and because attrition from the Year 1 to Year 2 study was differential based on treatment status (see Appendix C.8), this last specification is open to endogeneity concerns and therefore should not be interpreted causally. For the sake of transparency, we present it regardless, but with aforementioned caveat.

3.4 Individual level results

Take up

Take-up of the loan treatments was quite high. Of the 954 individuals in the Year 1 treatment group, 610 (64%) applied and qualified for the loan. In Year 2, 324 out of the 522 treated individuals (62%) qualified for and took up the loan. Unconditional loan sizes in the two treatment groups were 4,817 Ksh and 6,679 Ksh or about $57 and $79 USD), respectively. The average loan sizes conditional on take-up were 7,533 Ksh (or about $89 USD) for Year 1 and 10,548 Ksh (or $124) for Year 2. [22]

Relative to many other credit-market interventions in low-income settings in which documented take-up rates range from 1-10% of the surveyed population (Karlan, Morduch and Mullainathan, 2010), the 60-65% take-up rates of our loan product were extraordinarily high. This is perhaps not surprising given that our loan product was offered as a top-up for individuals who were already clients of an MFI. Nevertheless, OAF estimates that 20-30% of farmers in a given village in our study area enroll in OAF, which implies that even if *no* non-OAF farmers were to adopt the loan if offered it, population-wide take-up rates of our loan product would still exceed 10-20%.

Overall price increase

Before turning to the results regarding the impacts of the loan, it is worth noting the small average price increase that occurred during both study years, relative to what farmers (and traders) reported had occurred in the recent past, and relative to what was expected by farmers. As shown in the right panel of Figure 3.2, farmers had expected a doubling of prices, but prices only increased by 20-30% and peaked 2-3 months earlier than normal in

[22] Recall in Year 1 there were two version of the loan, one offered in October and the other in January. Of the 474 individuals in the 77 groups assigned to the October loan treatment (T1), 329 (69%) applied and qualified for the loan. For the January loan treatment (T2), 281 out of the 480 (59%) qualified for and took up the loan. Unconditional loan sizes in the two treatment groups were 5,294 Ksh and 4,345 Ksh (or about $62 and $51 USD) for T1 and T2, respectively, and we can reject at 99% confidence that the loan sizes were the same between groups. The average loan sizes conditional on take-up were 7,627Ksh (or about $90 USD) for T1 and 7,423Ksh (or $87) for T2, and in this case we cannot reject that conditional loan sizes were the same between groups.

both years. We currently do not know why this is – prices in larger surrounding markets were also flat – but we are currently conducting interviews with local traders to try to understand why the study years might have been different than the preceding period.

Primary effects of the loan offer

We begin by estimating treatment effects in the standard fashion, assuming that there could be within-randomization-unit spillovers (in our case, the group), but that there are no cross-group spillovers. In all tables and figures, we report results broken down by each year and pooled. As explained in Section 3.3, the Year 2 replication produced results that are quantitatively quite similar to the Year 1 results for most outcomes, and as such, we report in the text the pooled results, unless otherwise noted.

Tables 3.2-3.7 and Figure 3.3 and show the results of estimating Equations 3.1 and 3.2 on the pooled treatment indicator, either parametrically (in the table) or non-parametrically (in the figure). The top panels in Figure 3.3 show the means in the treatment group (broken down by year and then pooled, in the final panel) over time for our three main outcomes of interest (as estimated with fan regressions). The bottom panels show the difference in treatment minus control over time, with the 95% confidence interval calculated by bootstrapping the fan regression 1000 times.

Farmers responded to the intervention as anticipated. They held significantly more inventories for much of the year, on average about 25% more than the control group mean (Column 6 in Table 3.2). Inventory effects are remarkably similar across replication years.

Net revenues[23] were significantly lower immediately post harvest and significantly higher later in the year (Column 6 in Table 3.3 and middle panel of Figure 3.3). The net effect on revenues averaged across the year is positive in both years of the experiment, but is only significant in the Year 2 and the pooled data (see Columns 1, 3, and 5 in Table 3.3). Breaking down Year 1 results by the timing of loan suggest that the reason results in Year 1 are not significant is that the later loan, offered in January to half of the treatment group, was less effective than the October loan. Table C.1 presents results for the Year 1 loan, broken down by loan timing. We see in Column 5 that the October loan (T1) produced revenue effects that are more similar in magnitude (and now significant, at 5%) to those of the Year 2 loan (which was offered almost at the same time). The January loan (T2) had no significant effect on revenues. Appendix Section C.2 explores the effects of loan timing in greater detail. The total effect across the year can be calculated by adding up the coefficients in Column 6 of Table 3.3,which yields an estimate of 1616 Ksh, or about $19 at exchange rate at the time of the study. Given the unconditional average loan size of 5,476 Ksh in the pooled data, this is equivalent to a 30% return.

The final panel of Figure 3.3 and Table 3.4 present the consumption effects (as measured by logged total household consumption). While point estimates are positive in both years, they are insignificant when pooled (in Year 2, treatment is associated with a 7 percentage

[23] From which loan interest rates were subtracted for those who took out a loan.

point increase in consumption, significant at 10%, but in Year 1, estimated effects are only slightly greater than zero and are insignificant).

Tables 3.5 - 3.7 present outcomes on a few other outcomes of interest. Table 3.5 suggests that net sales of maize are a big larger in the treatment group (with the time trend of net sales as shown in Column 6 following the expected pattern, with lower net sales immediately after harvest and greater net sales later in the season). Table 3.6 and Table 3.7 present suggestive evidence that treated individual are able to purchase maize at lower prices (significant at 5% in the pooled data, 10% in Year 2, and insignificant in Year 1) and sell maize at a higher price (though the evidence on the latter point is less clear; results are insignificant in the pooled data and Year 1, and only significant at 10% in Year 2).

Secondary effects of the loan offer

Appendix Section C.4 presents outcomes on potential secondary outcomes of interest. We find no significant effects on profits earned from and hours worked at non-farm household-run businesses (Tables C.5, nor on C.6), wages earned from and hours worked in salaried employment (Tables C.7 and C.8). We also find no significant effects on schools fees paid (the primary expenditure that households say constrain them to sell their maize stocks early; see Table C.9). We do in Year 1 find a significant 0.07 point increase on a happiness index (an index for the following question: "Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)"). However, we find no significant increase in this measure in Year 2.

Long-run effects

Appendix Section C.5 explores the long-run follow-up effects of the loan, as measured in the Long-Run Follow-Up (LRFU) survey conducted November-December 2015, which measures outcomes one to two years after the completion of the intervention (for the Year 2 and Year 1 loan respectively). In this section, we primarily focus on the effects of each year of the study as estimated separately, because these results can be interpreted causally. In the tables in Appendix C.5, we also present the effects of the interaction of treatment in each year, but do not discuss these results here, because this specification cannot be interpreted causally (attrition from the sample between Years 1 and 2 was differential based on Year 1 treatment status).

We first explore outcomes for the 2014 long-rains and 2015 short-rains harvests, the two seasons immediately following the completion of the Year 2 study. We see no effects on harvest levels (Columns 1-3) or amount eaten (Columns 4-6) in Table C.11, nor do we see any effects on net sales (Column 1-3) or revenues (Column 4-6) in Table C.12. We also see no effects when broken down into amount and value sold (Table C.13) vs. purhcased (Table C.14), nor do we see any significant when exploring these outcomes separately by season (Tables C.15 - C.18).

We also explore effects for the 2015 long-rains harvest, one to two years following the conclusion of the loan. Again, we see no evidence of long-run effects on harvest amounts (Table C.19). Consistent with this, we see no increase in inputs used, either on labor or non-labor inputs (Table C.20). We therefore find no evidence that increased storage and revenues in one year crowds in other inputs and increases harvests in future years.

We also explore other outcomes for the 2015 year. Table C.21 explores long-run effects on maize eaten and food expenditures. We find no significant effects. Table C.22 presents impacts on overall household (log) consumption. Again, we find no significant effects. Table C.22 also explores the long-run effects on the happiness index (an index for the following question: "Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)"). We find an increase of 0.1 points on the index from the Year 1 treatment, which is around the same size as (in fact, a larger point estimate than) the immediate effects on happiness. However, we find no effect of the Year 2 treatment on the long-run happiness index (consistent with the lack of an immediate effect for this study year). Why one year had a greater effect on the happiness index is an issue deserving of further exploration.

Table C.23 explores long-run effects on educational expenditure and school attendance (as measured by the proportion of days the children in the household attended school in the last five days). We find no significant effects on either.

Table C.24 displays long-run effects on the hours spent on non-farm businesses owned by the household (Columns 1-3) and profits from these businesses (Columns 4-6). We see no significant effects. Table C.24 presents long-run impacts on hours work and wages at salaried employment positions. We find no effects on hours worked. The point estimate on wages are positive, but is only significant in Year 2.

Finally, Table presents long-run effects on prices from the period of November 2014 - August 2015 (one year after the loan was removed from the region). We no longer see any significant effects (and in fact the signs on the point estimates have switched).

In summary, we find little evidence of meaningful long-run impacts of the loan. This suggests that – at least in the context of this loan product – farmers are not able to save their way out of their credit constraint.

3.5 General equilibrium effects

The experiment was designed to quantify one particular potential general equilibrium effect: the effect of the loan intervention on local maize prices. Such effects appeared plausible for three reasons. First, OAF serves a substantial number of farmers in a given area. In "mature" areas where OAF has been working for a number of years (such as Webuye district where our experiment took place), typically 30% of all farmers sign up for OAF. This means that in high treatment density areas, where 80% of OAF groups were enrolled in the study and 2/3rds of these offered the loan, roughly 10% of the population of farmers took the loan. Second, focus groups had suggested take up of the loan would be quite high, and that farmers did not need to be told that they could make extra money by storing longer. Finally, while we lack long-term price data for local markets in the area, there is evidence that these markets are not well integrated. In particular, a handful of traders can be found in these markets on the main market day, and there is evidence that these traders are making substantial profits engaging in spatial arbitrage across these markets, often selling in markets they will later purchase from (and vice versa) (Bergquist, 2017). This suggests that these markets might be affected by local shifts in supply and demand.[24]

Market level effects

To understand the effect of our loan intervention on local maize prices, we identified 52 local market points spread throughout our study area that OAF staff indicated were where their clients typically bought and sold maize, and our enumerators tracked monthly maize prices at these market points. We then match these market points to the OAF sublocation in which they fall. "Sublocations" here are simply OAF administrative units that are well defined in terms of client composition (i.e. which OAF groups are in which sublocation), but less well defined in terms of their exact geographic boundaries. Given this, using GPS data on both the market location and the location of farmers in our study sample to calculate the "most likely" sublocation, based on the designated sublocation to which the majority of nearby farmers belong.[25].

We then utilize the sublocation-level randomization in treatment intensity to identify market-level effects of our intervention, estimating Equation 3.3 and clustering standard errors at the sublocation level. Regression results are shown in Table 3.8 and plotted non-parametrically in Figure 3.4. In each year, we explore the price changes from November (immediately following harvest) until August (the beginning of the subsequent season's harvest). In Figure 3.4, which presents the results pooling Year 1 and Year 2 of price data, we see prices in high-intensity areas start out about 3% higher in the immediate post-harvest months. As the season goes on, price in high density areas then begin to converge and even dip below those low density areas. Table 3.8 presents these results according to the empirical specifically outlined in Section 3.3. In line with the graphic results visible in Figure 3.4, here we see the interaction term on "Hi" treatment intensity is positive (and significant at 10%), while the interaction term between the monthly time trend and the high intensity dummy is negative (though not significant). The overall picture painted by the market price data is remarkably consistent with the individual-level results presented above. Larger inward shifts

[24]Other papers, such as Cunha, De Giorgi and Jayachandran (2011), find substantial effects of local supply shocks on local prices in settings (in this case, Mexico) where markets are likely much less isolated than ours.

[25]As a robustness check, in the Appendix Section C.3, we use an alternative matching method in which we ask OAF field staff which markets are in their sublocation. Results in Year 1 are robust to matching method, though in Year 2 first stage results are weaker and reduced formed results are less precise using this alternate method.

in supply early on caused prices to start higher in high-intensity areas, and prices equalize at about the time the treated individuals switch from being net buyers to net sellers.

Note that results are weaker in Year 2 than in Year 1 (though coefficients share the same, expected signs). This is likely because the assigned treatment intensity (which recall was kept constant from Year 1 to Year 2) is a weaker instrument for observed intensity of treatment in Year 2 compared to Year 1, due to administrative reshuffling of groups between the first and second year of the experiment. Table 3.12 quantifies this effect. The first stage of assigned intensity on observed in intensity in Year 2 is half that of Year 1. In Table 3.9, we correct the reduced form results for the differences in first stage effects by using the assigned intensity as an instrument for observed intensity separately in Year 1 (in Columns 1-2) and Year 2 (in Columns 3-4). We see IV effects that are remarkably similar across the two years (albeit less precise measured in Year 2, again due to the weaker first stage).

Appendix C.6 presents further robustness checks of these price effects.

Individual results with spillovers

We now revisit the individual results, re-estimating them to account for the variation in treatment density across sublocations. We note at the outset that while our experiment affected local market prices differentially in high- and low-treatment density areas, changes in treatment density could precipitate other spillovers beyond output price effects. For instance, sharing of maize or informal lending between households could also be affected by having a locally higher density of loan recipients; as an untreated household, your chance of knowing someone who got the loan is higher if you live in a high-treatment-density areas. Nevertheless, these spillovers could be positive or negative – e.g. we don't know *ex ante* whether our treatment would cause individuals to exit informal lending relationships or to expand them, or whether it would allow them to reduce their maize transfers or allow them to give out more maize to untreated households. We attempt to clarify the sign and magnitude of these potential spillovers in what follows.

Tables 3.10 - 3.17 and Figure 3.7 show how our three main outcomes respond in high versus low density areas for treated and control individuals. Inventory treatment effects do not significantly differ as a function of treatment intensity for the pooled treatment. Effects on net revenues, however, paint a different picture. Treatment effects in low intensity areas are much larger than what was estimated earlier. In contrast, revenue effects for treated individuals in high intensity areas are lower (and in fact are statistically indistinguishable from zero in Column 3 of Table 3.11). This suggests that there is something about higher treatment density that erodes the effect of the loan on profitability. There is also some evidence that net revenues were higher in high-intensity control group relative to the low intensity control group (see middle panel of Figure 3.7 and the estimate on the Hi dummy in Column 3 of Column 3 of Table 3.11; though it should be noted that the effect is insignificant and, for Year 2, even goes in the wrong direction (see Column 2 of Table 3.11). Columns 4-6 of Table 3.13 present the instrumented version of these results. After accounting for the weaker first stage in Year 2, we see remarkably similar revenue effects of treatment across the two years (2,645 in Year 1 and 2,345 in Year 2) and of treatment interacted with observed treatment intensity (-9,111 in Year 1 and -10,684 in Year 2) (see Columns 4 and 5).

Effects on consumption, as with earlier estimates, remain quite noisy.[26]

Could these differential net revenue effects have come through price spillovers alone? Here we explore some more prosaic alternative explanations. First, note that covariates were balanced at baseline in Year 1 between high- and low-intensity areas (Table C.27), so we can rule out simple concerns of imbalance.

We do find some imbalances in loan take-up by intensity (see Table 3.18). In high intensity areas, loan take-up is 7 percentage points lower than in low areas (significant at 1%) overall (Row 3), though interestingly, this pattern reverses from Year 1 (when loan take-up is 13 percentage points lower in high intensity areas) to Year 2 (when loan take-up is 6 percentage points higher in high intensity areas).[27] This differential take-up could matter for our treatment effects because we estimate the Intent-to-treat, and given a constant treatment-effect-on-the-treated, ITT estimates should be mechanically closer to zero in cases where take-up is lower. One might worry that, in particular in Year 1 when take-up is lower in the high intensity areas, this explains why revenue effects are also lower in high intensity areas. Two factors argue against this concern. First, the difference appears too small to explain our results fully. If there were no other spillovers, and treatment-on-treated effects were the same in high and low intensity areas, then ITT estimates in the high intensity ares should be 83% as large (0.63/0.76). However, point estimates on revenue treatment effects in Year 1 are roughly *zero* in the high-intensity areas (compared to 1,060 in low-intensity areas), a much bigger gap that could be explained by differential take-up. Second, and moreover, in Year 2, the differential take-up pattern switches; in this year, take-up is *higher* in high-intensity areas. If take-up were driving these results, we should see that a switch in the take-up patterns by intensity results in a switch in the revenue effects by intensity. However, we consistently across Years 1 and 2 see that revenue effects are greater among low-intensity areas. Take-up is therefor unlikely to be driving results.

We do additionally see some differences in loan size by intensity in Year 2. In this year of the experiment, loans were larger in high intensity areas. However, this should have driven *greater* revenue effects in high intensity areas, rather than the lower effects that we find. We therefore believe it is unlikely that differential take-up or loan size are driving these results.

Overall, then, the individual-level spillover results are perhaps most consistent with spillovers through market prices. We conclude this section by noting that, had we just run the experiment at our high treatment density, we would have found results very similar

[26]Interestingly, they are strongly positive for treated individuals in the high-intensity areas in Year 2. However, because there is no clear pattern across years and because the other coefficients are so imprecisely measured, we avoid speculating or over-interpreting this figure.

[27]The Year 1 results may be the result of repayment incentives faced by OAF field staff: our loan intervention represented a substantial increase in the total OAF credit outlay in high-intensity areas, and given contract incentives for OAF field staff that reward a high repayment rate for clients in their purview, these field officers might have more carefully screened potential adopters. We are still exploring why the Year 2 results would have switched.

to what has been found in existing microcredit literature: a significant effect of improved credit access on inventories, but zero effect on revenues. While our rural setting is one in which certain types of spillovers (e.g. through prices) might be more important relative to the more urban settings that typify the existing microcredit experiments, our results do suggest that "headline" estimates of microcredit's impacts could be substantially shaped by the saturation at which the experiment is run.

3.6 Conclusion

We study the effect of offering Kenyan maize farmers a cash loan at harvest. The timing of this loan is motivated by two facts: the large observed average increase in maize prices between the post harvest season and the lean season six to nine months later, and the inability of most poor farmers appear to successfully arbitrage these prices due to a range of "non-discretionary" consumption expenditures they must make immediately after harvest. Instead of putting maize in storage and selling when the price is higher, farmers are observed to sell much of it immediately, sacrificing potential profits.

We show that access to credit at harvest "frees up" farmers to use storage to arbitrage these prices. Farmers offered the loan shift maize purchases into the period of low prices, put more maize in storage, and sell maize at higher prices later in the season, increasing farm profits. Using experimentally-induced variation in the density of treatment farmers across locations, we document that this change in storage and marketing behavior aggregated across treatment farmers also affects local maize prices: post harvest prices are significantly higher in high-density areas, consistent with more supply having been taken off the market in that period, and are lower later in the season (but not significantly so). These general equilibrium effects feed back to our profitability estimates, with farmers in low-density areas – where price differentials were higher and thus arbitrage opportunities greater – differentially benefiting.

Our findings make a number of contributions. First, our results are some of the first experimental results to find a positive and significant effect of microcredit on the profits of microenterprises (farms in our case), and the first experimental study to directly account for general equilibrium effects in this literature. While we cannot claim that these two facts are more generally related, it is the case in our particular setting that failing to account for these GE effects substantially alters the conclusions drawn about the average benefits of improved credit access. This suggests that explicit attention to GE effects in future evaluations of credit market interventions could be warranted.

Second, we show how the absence of financial intermediation can be doubly painful for poor households in rural areas. Lack of access to formal credit causes households to turn to much more expensive ways of moving consumption around in time, and aggregated across households this behavior generates a broad scale price phenomenon that further lowers farm income and increases what these households must pay for food. Our results suggest that in this setting, expanding access to affordable credit could reduce this price variability and thus have benefits for recipient and non-recipient households alike.

What our results do not address is why larger actors – e.g. large-scale private traders – have not stepped in to bid away these arbitrage opportunities. It is perhaps surprising that traders are not more active in inter-temporal arbitrage, given that they are less credit constrained than individual farmers. To answer this question (and as part of the pilot work conducted prior to the experiments run in Chapters 1 and 2 of this dissertation), panel surveys were conducted with 359 maize traders in the region.

Interviewed traders on average report that they could make a 30% return buying maize at harvest, storing, and selling later in the season after prices rise. However, few traders exploit this apparent arbitrage opportunity. Among traders surveyed in the panel study, only 8% stored for longer than 2 months; those who do store for this period typically only store for 3 months, substantially less than the eight-month period from trough to peak prices.

One of the most commonly cited explanations for why traders do not engage in inter-temporal arbitrage is that they are able to earn even higher total profits by engaging in spatial arbitrage (perhaps because imperfect competition among traders is at play (Bergquist, 2017)).[28] Rates of return calculated from self-reported profits confirm this, suggesting that spatial arbitrage opportunities to move maize from surplus to deficit markets trump the temporal arbitrage opportunities offered by storage.[29]

This suggests that policymakers may need consider the interaction of inefficiencies in spatial and temporal arbitrage, as opportunities to earn rents in one dimension may move intermediaries out of arbitrage in the other.

[28] 60% of traders cited this as the reason they do not store.

[29] While traders report at 30% return on storing a bag of maize for six months, they report being able to make a 12% return on spatial arbitrage in less than a week, for annualized returns of much larger than that for storage (if traders conduct spatial arbitrage every week for six months, instead of storing, this would yield a return of about 300%). Further, structural estimates from Bergquist (2017) suggest that the self-reported mark-ups on spatial arbitrage may be severely under-reported and that actual mark-ups on spatial arbitrage may be much greater than 12%.

Tables and Figures

Figure 3.1: **M**onthly average maize prices, shown at East African sites for which long-term data exist, 1994-2011. Data are from the Regional Agricultural Trade Intelligence Network, and prices are normalized such that the minimum monthly price = 100. Our study site in western Kenya is shown in green, and the blue squares represent an independent estimate of the months of the main harvest season in the given location. Price fluctuations for maize (corn) in the US are shown in the lower left for comparison

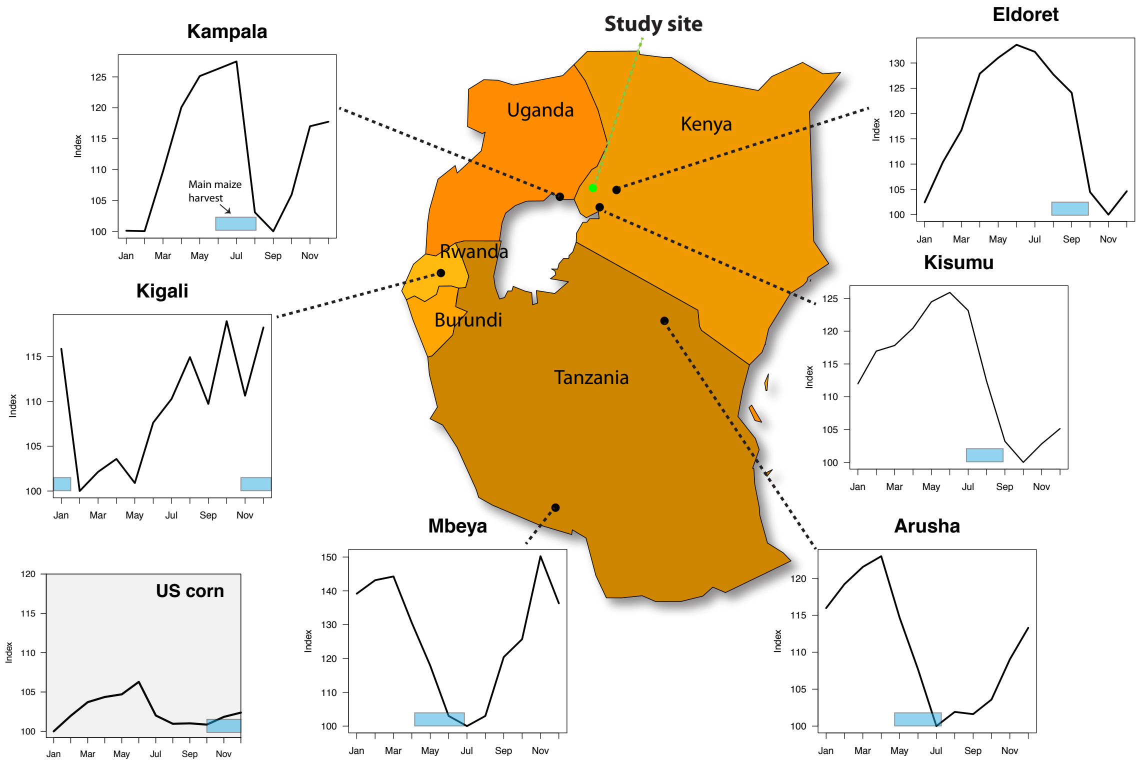

Figure 3.2: **M**aize prices in local markets. **L**eft panel: farmer-reported average monthly maize prices for purchase and sales over 2007-2012, averaged over all farmers in our sample. Prices are in Kenyan shillings per goro goro (2.2kg). **R**ight panel: farmers expectations of sales prices over the Sept2012-Aug2013 period, as reported in August2012 (solid red line), and actual observed sales prices in local markets over the same period (dotted line).

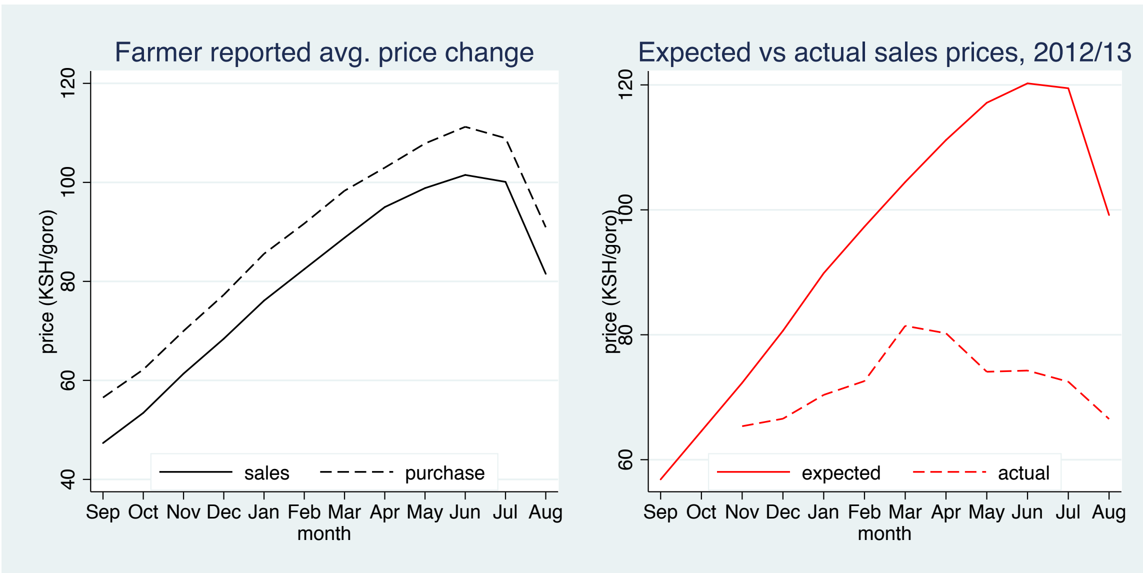

Figure 3.3: **P**ooled Treatment effects. The top row of plots shows how average inventories, net revenues, and log household consumption evolve from December to August in Y1 and Y2 (pooled) in the treatment group versus the control group, as estimated with fan regressions. The bottom row shows the difference between the treatment and control, with the bootstrapped 95% confidence interval shown in grey (100 replications drawing groups with replacement).

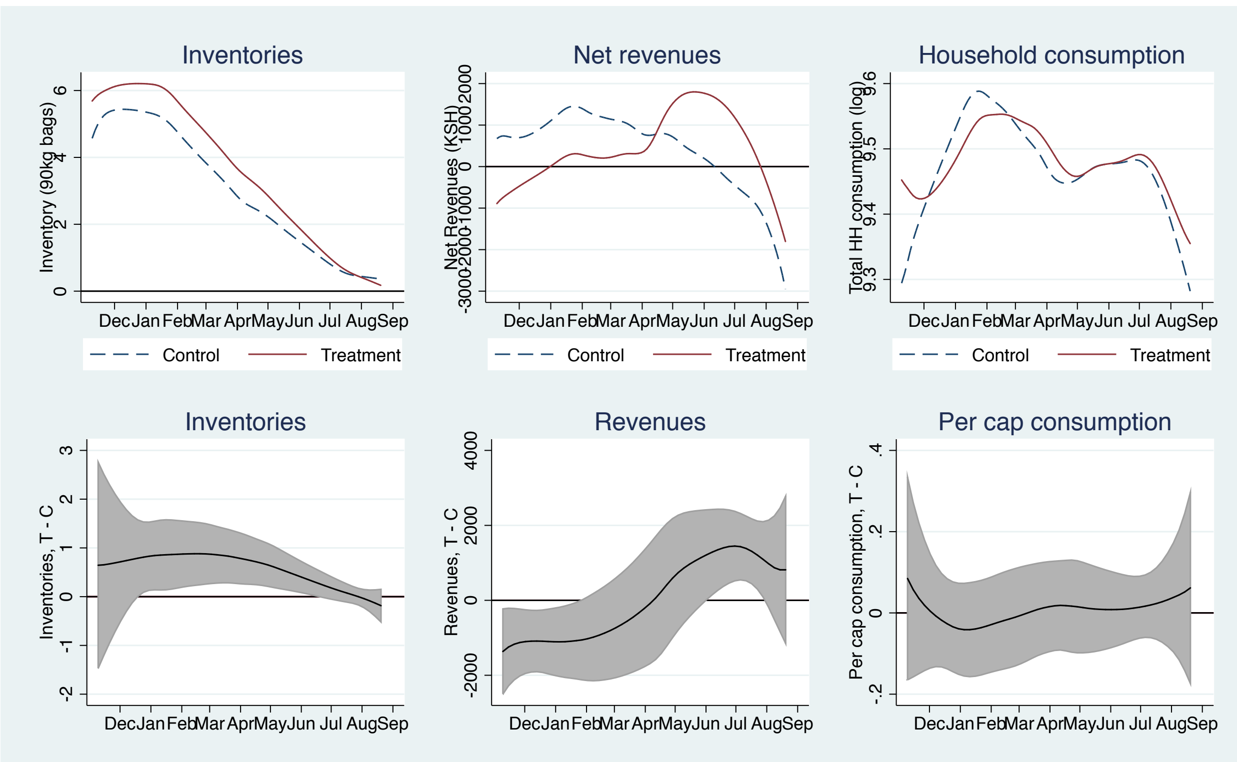

Figure 3.4: **P**ooled Market prices for maize as a function of local treatment intensity (nearest). The left panel shows the average sales price in markets in high-intensity areas (solid line) versus in low-intensity areas (dashed line) over the study period. The right panel shows the average difference in log price between high- and low-intensity areas over time, with the bootstrapped 95% confidence interval shown in grey. Markets matched to treatment intensity using location data on farmers and markets.

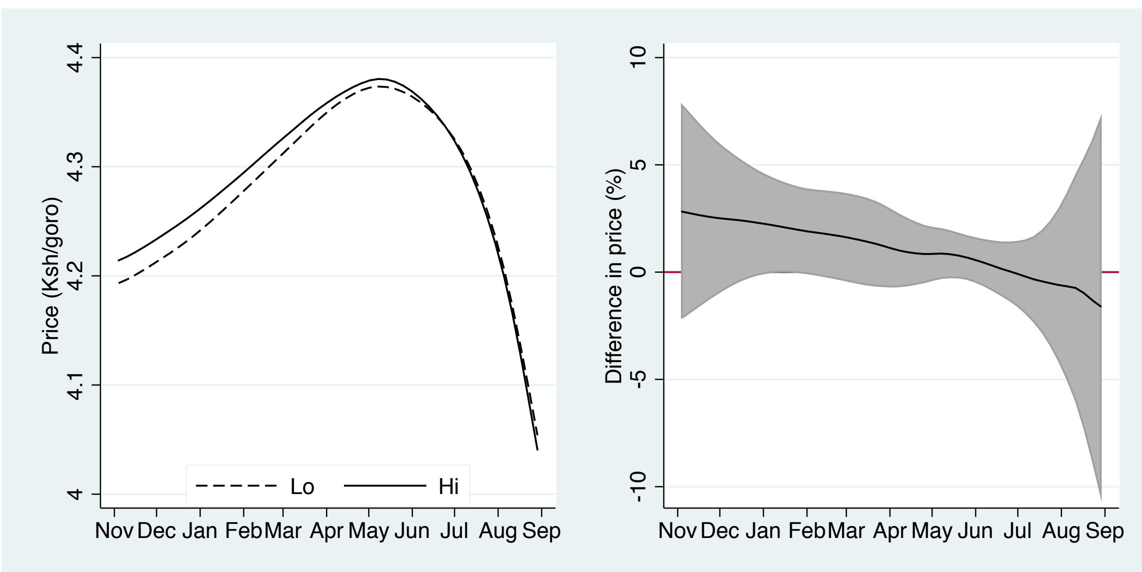

Figure 3.5: **Y**1 Market prices for maize as a function of local treatment intensity (nearest). The left panel shows the average sales price in markets in high-intensity areas (solid line) versus in low-intensity areas (dashed line) over the study period. The right panel shows the average difference in log price between high- and low-intensity areas over time, with the bootstrapped 95% confidence interval shown in grey. Markets matched to treatment intensity using location data on farmers and markets.

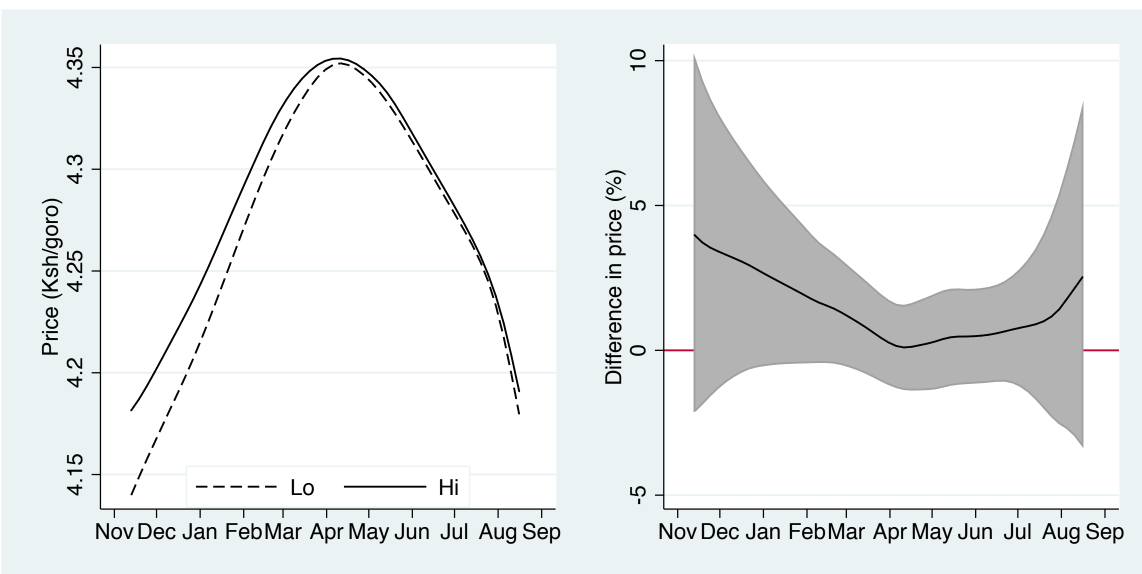

Figure 3.6: **Y**2 Market prices for maize as a function of local treatment intensity (nearest). The left panel shows the average sales price in markets in high-intensity areas (solid line) versus in low-intensity areas (dashed line) over the study period. The right panel shows the average difference in log price between high- and low-intensity areas over time, with the bootstrapped 95% confidence interval shown in grey. Markets matched to treatment intensity using location data on farmers and markets.

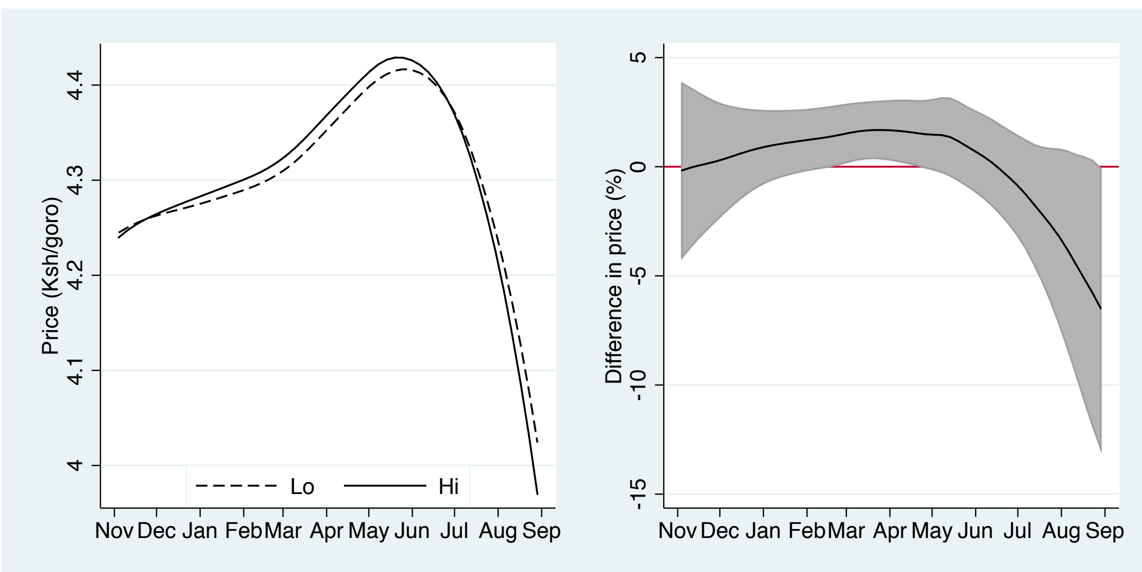

Figure 3.7: **P**ooled Treatment effects by treatment intensity. Average inventories, net revenues, and log HH consumption over the study period in the treatment group versus the control group, split apart by high intensity areas (orange lines) and low-intensity areas (black lines).

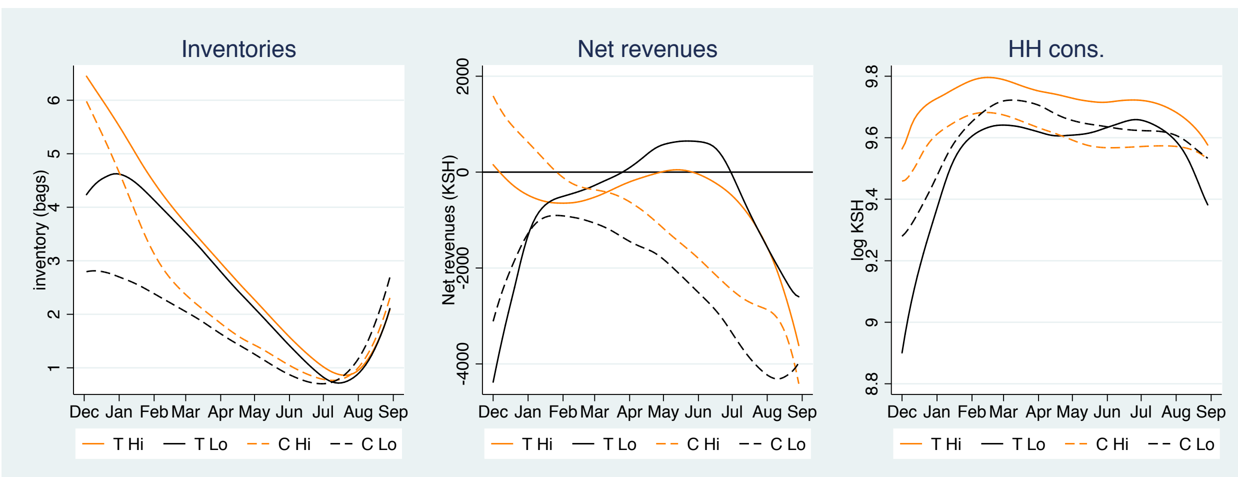

Table 3.1: **S**ummary statistics and balance among baseline covariates. Balance table for the Y1 study (restricted to the Y1 sample, for which we have baseline characteristics. The first two columns give the means in each treatment arm. The 3rd column gives the total number of observations across the two groups. The last two columns give differences in means normalized by the Control sd, with the corresponding p-value on the test of equality.

Baseline characteristic	C	Treat12	Obs	C - Treat12	
				sd	*p-val*
Male	0.33	0.30	1,589	0.08	0.11
Number of adults	3.20	3.00	1,510	0.09	0.06
Kids in school	3.07	3.00	1,589	0.04	0.46
Finished primary	0.77	0.72	1,490	0.13	0.02
Finished secondary	0.27	0.25	1,490	0.04	0.46
Total cropland (acres)	2.40	2.44	1,512	-0.01	0.79
Number of rooms in hhold	3.25	3.07	1,511	0.05	0.17
Total school fees (1000 Ksh)	29.81	27.24	1,589	0.06	0.18
Average monthly cons (Ksh)	15,371.38	14,970.86	1,437	0.03	0.55
Avg monthly cons./cap (log Ksh)	7.96	7.97	1,434	-0.02	0.72
Total cash savings (KSH)	8,021.50	5,157.40	1,572	0.09	0.01
Total cash savings (trim)	5,389.84	4,731.62	1,572	0.05	0.33
Has bank savings acct	0.43	0.42	1,589	0.01	0.82
Taken bank loan	0.08	0.08	1,589	0.02	0.73
Taken informal loan	0.25	0.24	1,589	0.01	0.84
Liquid wealth	97,280.92	93,878.93	1,491	0.03	0.55
Off-farm wages (Ksh)	3,797.48	3,916.82	1,589	-0.01	0.85
Business profit (Ksh)	1,801.69	2,302.59	1,589	-0.08	0.32
Avg %Δ price Sep-Jun	133.18	133.49	1,504	-0.00	0.94
Expect 2011 LR harvest (bags)	9.03	9.36	1,511	-0.02	0.67
Net revenue 2011	-4,088.62	-3,303.69	1,428	-0.03	0.75
Net seller 2011	0.30	0.32	1,428	-0.05	0.39
Autarkic 2011	0.06	0.07	1,589	-0.03	0.51
% maize lost 2011	0.01	0.02	1,428	-0.03	0.57
2012 LR harvest (bags)	11.03	11.18	1,484	-0.02	0.74
Calculated interest correctly	0.73	0.71	1,580	0.03	0.50
Digit span recall	4.58	4.57	1,504	0.01	0.89
Maize giver	0.26	0.26	1,589	0.00	0.99
Delta	0.13	0.13	1,512	-0.04	0.43

"Liquid wealth" is the sum of cash savings and assets that could be easily sold (e.g. livestock). Off-farm wages and business profit refer to values over the previous month. Net revenue, net seller, and autarkic refer to the household's maize marketing position. "Maize giver" is whether the household reported giving away more maize in gifts than it received over the previous 3 months. "Delta" is the percent of allocations to the earlier period in a time preference elicitation.

Table 3.2: **I**nventory Effects, Individual Level. Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	0.52***		0.50***		0.53***	
	(0.16)		(0.14)		(0.12)	
Treat - R1		0.82***		1.21***		1.03***
		(0.31)		(0.24)		(0.20)
Treat - R2		0.71***		0.24		0.52***
		(0.19)		(0.15)		(0.12)
Treat - R3		0.06		0.04		0.07
		(0.07)		(0.37)		(0.19)
Observations	3836	3836	2944	2944	6780	6780
Mean DV	2.67	2.67	1.68	1.68	2.16	2.16
R squared	0.35	0.35	0.18	0.19	0.29	0.30

Table 3.3: **N**et Revenue Effects, Individual Level. Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	279.78		800.24**		524.66**	
	(292.16)		(330.63)		(220.25)	
Treat - R1		-1146.56***		-23.71		-608.68**
		(325.13)		(478.41)		(285.70)
Treat - R2		534.85		1917.28***		1170.71***
		(485.80)		(532.81)		(359.84)
Treat - R3		1371.95***		520.76		985.79***
		(436.12)		(403.27)		(302.09)
Observations	3795	3795	2935	2935	6730	6730
Mean DV	334.41	334.41	-3434.38	-3434.38	-1616.12	-1616.12
R squared	0.01	0.01	0.04	0.05	0.09	0.09

Table 3.4: **H**H Consumption (log) Effects, Individual Level. Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	0.00		0.07*		0.04	
	(0.03)		(0.04)		(0.03)	
Treat - R1		-0.04		0.07		0.01
		(0.05)		(0.05)		(0.03)
Treat - R2		0.02		0.08*		0.05
		(0.04)		(0.05)		(0.03)
Treat - R3		0.03		0.06		0.04
		(0.05)		(0.05)		(0.03)
Observations	3792	3792	2944	2944	6736	6736
Mean DV	9.48	9.48	9.61	9.61	9.55	9.55
R squared	0.00	0.00	0.01	0.01	0.02	0.02

Table 3.5: **N**et Sales Effects, Individual Level. Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	0.07		0.19**		0.12*	
	(0.10)		(0.07)		(0.06)	
Treat - R1		-0.47***		0.08		-0.26**
		(0.13)		(0.16)		(0.10)
Treat - R2		0.16		0.44***		0.27***
		(0.15)		(0.12)		(0.10)
Treat - R3		0.48***		0.04		0.29***
		(0.14)		(0.11)		(0.09)
Observations	3820	3820	2288	2288	6108	6108
Mean DV	0.18	0.18	-1.55	-1.55	-0.62	-0.62
R squared	0.01	0.01	0.01	0.01	0.16	0.16

Table 3.6: **P**urchase Price Effects, Individual Level. Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	-19.22		-32.63*		-26.12**	
	(18.26)		(16.63)		(12.44)	
Treat - R1		-32.77		-18.23		-22.35
		(45.82)		(30.36)		(25.14)
Treat - R2		-16.77		-19.04		-16.89
		(30.73)		(23.90)		(19.50)
Treat - R3		-15.22		-56.57*		-36.07**
		(18.53)		(30.18)		(17.99)
Observations	1908	1908	2282	2282	4190	4190
Mean DV	2982.26	2982.26	3310.28	3310.28	3193.11	3193.11
R squared	0.36	0.36	0.33	0.33	0.45	0.45

Table 3.7: **S**ales Price Effects, Individual Level. Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	12.85		52.75*		24.53	
	(22.14)		(31.71)		(18.19)	
Treat - R1		-32.51		57.89		8.35
		(43.10)		(36.69)		(29.19)
Treat - R2		37.86		78.99		46.70*
		(30.07)		(51.35)		(25.59)
Treat - R3		24.89		-32.66		13.96
		(31.51)		(125.25)		(35.03)
Observations	1424	1424	636	636	2060	2060
Mean DV	2830.29	2830.29	3024.90	3024.90	2899.93	2899.93
R squared	0.44	0.45	0.20	0.20	0.41	0.41

Table 3.8: **M**arket prices for maize as a function of local treatment intensity (nearest).

	Y1		Y2		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
Hi	2.81*	2.80*	0.99	0.97	2.00*	2.00*
	(1.41)	(1.46)	(0.93)	(0.89)	(1.01)	(1.07)
Time	0.78***	0.78***	0.88***	0.88***	0.86***	0.86***
	(0.24)	(0.24)	(0.20)	(0.20)	(0.20)	(0.20)
Hi Intensity * Time	-0.39	-0.39	-0.15	-0.15	-0.28	-0.28
	(0.27)	(0.27)	(0.24)	(0.24)	(0.23)	(0.23)
Observations	491	491	454	454	945	945
Mean of Dep Var	61.87	61.87	61.87	61.87	61.87	61.87
R squared	0.08	0.08	0.06	0.06	0.06	0.07
Controls	No	Yes	No	Yes	No	Yes

Data are for November through August in Y1 and Y2. "Hi intensity" is a dummy for a sublocation randomly assigned a high number of treatment groups and "Time" is month number (beginning in November at 0 in each year). Standard errors are clustered at the sublocation level. Markets are matched to sublocations using location data on farmers and markets. Controls are the distance to the nearest road.

Table 3.9: **M**arket prices for maize as a function of local treatment intensity. IV regression (nearest).

	Y1		Y2		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
Observed Intensity	15.55**	15.66**	14.01	12.67	15.72**	16.09**
	(6.83)	(7.04)	(15.03)	(13.55)	(7.29)	(7.83)
Observed Intensity * Time	-2.34	-2.33	-1.84	-1.85	-2.21	-2.19
	(1.63)	(1.63)	(2.97)	(2.90)	(1.70)	(1.70)
Time	0.99**	0.99**	0.98**	0.98**	1.02***	1.01***
	(0.43)	(0.43)	(0.44)	(0.43)	(0.36)	(0.37)
Observations	434	434	407	407	841	841
Mean of Dep Var	66.66	66.66	66.66	66.66	66.66	66.66
R squared	0.06	0.06	0.05	0.05	0.04	0.04
Controls	No	Yes	No	Yes	No	Yes

Data are for November through August in Y1 and Y2. Assigned hi low treatment intensity is used as instruments for observed intensities, where the observed intensity is the number of treated farmers divided by the number of OAF farmers in that sublocation. "Time" is month number (beginning in November at 0 in each year). Standard errors are clustered at the sublocation level. Markets are matched to sublocations using location data on farmers and markets. Controls are the distance to the nearest road.

Table 3.10: **I**nventory Effects, Accounting for Treatment Intensity. Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1)	(2)	(3)
	Y1	Y2	Pooled
Treat	0.76***	0.55***	0.74***
	(0.19)	(0.18)	(0.15)
Hi	0.12	-0.03	0.02
	(0.36)	(0.22)	(0.24)
Treat*Hi	-0.33	-0.07	-0.29
	(0.23)	(0.25)	(0.19)
Observations	3836	2944	6780
Mean DV	2.74	1.38	2.04
R squared	0.35	0.18	0.29
p-val P+PH=0	0.01	0.02	0.01

Table 3.11: **N**et Revenue Effects, Accounting for Treatment Intensity. Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1) Y1	(2) Y2	(3) Pooled
Treat	1059.60**	1193.77	1101.39**
	(437.73)	(685.05)	(430.09)
Hi	533.90	-152.60	164.94
	(551.18)	(558.95)	(479.68)
Treat*Hi	-1114.63*	-555.21	-816.77
	(535.59)	(804.86)	(520.04)
Observations	3795	2935	6730
Mean DV	-253.51	-3620.40	-1980.02
R squared	0.01	0.04	0.09
p-val P+PH=0	0.86	0.15	0.41

Table 3.12: **F**irst Stage. First stage effect of assigned treatment intensity (0.4 or 0.8) on observed treatment intensity.

	(1) Obs Intens	(2) Obs Intens	(3) Obs Intens
Asg. Intens	0.38***	0.19**	0.31***
	(0.05)	(0.07)	(0.05)
Observations	3741	2251	5992
Mean DV	0.14	0.09	0.12
R squared	0.79	0.29	0.42
Type	FS	FS	FS
Sample	Y1	Y2	Pooled

Table 3.13: **R**F and IV Effects on Revenues. Reduced form effects of assigned treatment intensity on revenues. IV effect of observed treatment intensity (instrumented by assigned treatment intensity) on revenues.

	(1)	(2)	(3)	(4)	(5)	(6)
	Rev	Rev	Rev	Rev	Rev	Rev
Treat	2124.69**	1724.31	1881.86**	2620.63**	2345.75	2411.18**
	(903.20)	(1400.93)	(871.65)	(907.10)	(1844.30)	(947.27)
Asg. Intens	1305.10	-373.03	403.18			
	(1347.33)	(1366.32)	(1172.56)			
Treat x Asg. Intens	-2724.65*	-1357.19	-1996.55			
	(1309.23)	(1967.45)	(1271.20)			
Obs. Intens				5780.50	-1164.11	2998.64
				(3597.85)	(9982.16)	(5262.57)
Treat x Obs. Intens				-9087.41**	-10684.71	-9163.77
				(3867.34)	(13483.29)	(5188.30)
Observations	3795	2935	6730	3129	2220	5349
Mean DV	334.41	-3434.38	-1616.12	334.41	-3434.38	-1616.12
R squared	0.01	0.04	0.09	0.01	0.03	0.08
Type	RF	RF	RF	IV	IV	IV
Sample	Y1	Y2	Pooled	Y1	Y2	Pooled

Table 3.14: **H**H Consumption (log), Accounting for Treatment Intensity. Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1) Y1	(2) Y2	(3) Pooled
Treat	0.01	-0.05	-0.01
	(0.04)	(0.04)	(0.02)
Hi	-0.00	-0.08	-0.05
	(0.05)	(0.05)	(0.04)
Treat*Hi	-0.01	0.17***	0.07*
	(0.05)	(0.06)	(0.04)
Observations	3792	2944	6736
Mean DV	9.47	9.65	9.56
R squared	0.00	0.02	0.03
p-val P+PH=0	0.97	0.01	0.08

Table 3.15: **N**et Sales Effects, Accounting for Treatment Intensity. Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1) Y1	(2) Y2	(3) Pooled
Treat	0.39**	0.34***	0.38***
	(0.14)	(0.09)	(0.09)
Hi	0.21	0.29**	0.24*
	(0.17)	(0.12)	(0.12)
Treat*Hi	-0.45**	-0.21	-0.38***
	(0.18)	(0.13)	(0.12)
Observations	3820	2288	6108
Mean DV	-0.05	-1.80	-0.84
R squared	0.01	0.01	0.16
p-val P+PH=0	0.54	0.23	1.00

Table 3.16: **P**urchase Price Effects, Accounting for Treatment Intensity. Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1) Y1	(2) Y2	(3) Pooled
Treat	-21.59	-30.42	-27.38
	(25.35)	(25.65)	(20.97)
Hi	35.88	-21.23	-0.73
	(25.80)	(34.38)	(29.26)
Treat*Hi	-0.55	-2.55	1.78
	(31.81)	(37.15)	(28.50)
Observations	1908	2282	4190
Mean DV	2946.19	3291.45	3160.84
R squared	0.36	0.33	0.45
p-val P+PH=0	0.28	0.23	0.18

Table 3.17: **S**ales Price Effects, Accounting for Treatment Intensity. Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1) Y1	(2) Y2	(3) Pooled
Treat	44.58	79.87*	51.39
	(44.15)	(45.52)	(30.04)
Hi	83.90*	46.06	71.54*
	(46.76)	(46.58)	(36.61)
Treat*Hi	-44.44	-38.15	-37.61
	(46.94)	(53.60)	(32.14)
Observations	1424	636	2060
Mean DV	2760.62	3005.41	2849.41
R squared	0.45	0.20	0.41
p-val P+PH=0	0.99	0.12	0.30

Table 3.18: **L**oan Take-up and Size by Treatment Intensity.

	Loan Take-up					Loan Size				
	Low Mean	High Mean	N Obs	Diff SD	Diff p-val	Low Mean	High Mean	N Obs	Diff SD	Diff p-val
Year 1	0.76	0.63	2,703	0.30	0.00	7,426.27	7,576.19	1,804	-0.06	0.24
Year 2	0.59	0.65	1,354	-0.12	0.04	5,240.49	7,263.94	1,559	-0.38	0.00
Pooled	0.71	0.64	4,057	0.15	0.00	6,484.09	7,426.91	3,363	-0.23	0.00

Bibliography

Acemoglu, Daron. 2010. "Theory, general equilibrium and political economy in development economics." *Journal of Economic Perspectives*, 24(3): 17–32.

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

Aker, Jenny C. 2012. "Rainfall shocks, markets and food crises: the effect of drought on grain markets in Niger." *Center for Global Development, working paper.*

Angelucci, Manuela, Dean Karlan, and Jonathan Zinman. 2013. "Win some lose some? Evidence from a randomized microcredit program placement experiment by Compartamos Banco." National Bureau of Economic Research.

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

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

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

Attanasio, Orazio, Britta Augsburg, Ralph De Haas, Emla Fitzsimons, and Heike Harmgart. 2011. "Group lending or individual lending? Evidence from a randomised field experiment in Mongolia."

Baland, Jean-Marie, Catherine Guirkinger, and Charlotte Mali. 2011. "Pretending to be poor: Borrowing to escape forced solidarity in Cameroon." *Economic Development and Cultural Change*, 60(1): 1–16.

Banerjee, Abhijit V, and Andrew F Newman. 1993. "Occupational choice and the process of development." *Journal of political economy*, 274–298.

Banerjee, Abhijit V, and Esther Duflo. 2010. "Giving credit where it is due." *The Journal of Economic Perspectives*, 24(3): 61–79.

Banerjee, Abhijit Vinayak. 2013. "Microcredit Under the Microscope: What Have We Learned in the Past Two Decades, and What Do We Need to Know?" *Annual Review of Economics*, , (0).

Banerjee, A.V., E. Duflo, R. Glennerster, and C. Kinnan. 2013. "The Miracle of Microfinance?: Evidence from a Randomized Evaluation." *working paper, MIT*.

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

Barrett, C. 2008. "Displaced distortions: Financial market failures and seemingly inefficient resource allocation in low-income rural communities." *working paper, Cornell.*

Basu, Karna, and Maisy Wong. 2012. "Evaluating Seasonal Food Security Programs in East Indonesia." *working paper.*

Berge, Lars Ivar, Kjetil Bjorvatn, and Bertil Tungodden. 2011. "Human and financial capital for microenterprise development: Evidence from a field and lab experiment." *NHH Dept. of Economics Discussion Paper*, , (1).

Berge, Lars Ivar Oppendal, Kjetil Bjorvatn, Simon Galle, Edward Miguel, Daniel Posner, Bertil Tungodden, and Kelly Zang. 2016. "How Strong are Ethnic Preferences?" *Working Paper.*

Bergquist, Lauren. 2017. "Pass-through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya."

Blattman, Christopher, Nathan Fiala, and Sebastian Martinez. 2013. "Credit Constraints, Occupational Choice, and the Process of Development: Long Run Evidence from Cash Transfers in Uganda." *working paper.*

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2013. "Does management matter? Evidence from India." *The Quarterly Journal of Economics*, 128(1): 1–51.

Bruhn, Miriam, and David McKenzie. 2009. "In Pursuit of Balance: Randomization in Practice in Development Field Experiments." *American Economic Journal: Applied Economics*, 200–232.

Bruhn, Miriam, Dean S Karlan, and Antoinette Schoar. 2012. "The impact of consulting services on small and medium enterprises: Evidence from a randomized trial in mexico." *Yale University Economic Growth Center Discussion Paper*, , (1010).

Brune, L., X. Giné, J. Goldberg, and D. Yang. 2011. "Commitments to save: A field experiment in rural malawi." *University of Michigan, May (mimeograph).*

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

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

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller. 2008. “Bootstrap-based improvements for inference with clustered errors.” *The Review of Economics and Statistics*, 90(3): 414–427.

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

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

Casey, Katherine, Rachel Glennerster, and Edward Miguel. 2012. “Reshaping Institutions: Evidence on Aid Impacts Using a Preanalysis Plan*.” *The Quarterly Journal of Economics*, 127(4): 1755–1812.

Cohen, Jessica, and Pascaline Dupas. 2010. “Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment.” *Quarterly Journal of Economics*.

Crepon, B., F. Devoto, E. Duflo, and W. Pariente. 2011. “Impact of microcredit in rural areas of Morocco: Evidence from a Randomized Evaluation.” *working paper, MIT*.

Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran. 2011. “The price effects of cash versus in-kind transfers.” National Bureau of Economic Research.

De Mel, Suresh, David McKenzie, and Christopher Woodruff. 2008. “Returns to capital in microenterprises: evidence from a field experiment.” *The Quarterly Journal of Economics*, 123(4): 1329–1372.

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

Dupas, P., and J. Robinson. 2013. “Why Don’t the Poor Save More? Evidence from Health Savings Experiments.” *American Economic Review, forthcoming*.

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

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

Fafchamps, Marcel, and Eleni Gabre-Madhin. 2006. “Agricultural Markets in Benin and Malawi.” *African Journal of Agricultural and Resource Economics*, 1(1).

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

Fafchamps, Marcel, David McKenzie, Simon Quinn, and Christopher Woodruff. 2013. "Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana." *Journal of Development Economics*.

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

Field, Erica, Rohini Pande, John Papp, and Natalia Rigol. 2012. "Does the Classic Microfinance Model Discourage Entrepreneurship Among the Poor? Experimental Evidence from India." *American Economic Review*.

Galor, Oded, and Joseph Zeira. 1993. "Income distribution and macroeconomics." *The review of economic studies*, 60(1): 35–52.

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

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

Kaboski, Joseph P, and Robert M Townsend. 2012. "The impact of credit on village economies." *American economic journal. Applied economics*, 4(2): 98.

Karlan, Dean, and Jonathan Morduch. 2009. "Access to Finance." *Handbook of Development Economics, Volume 5*, , (Chapter 2).

Karlan, Dean, and Jonathan Zinman. 2011. "Microcredit in theory and practice: using randomized credit scoring for impact evaluation." *Science*, 332(6035): 1278–1284.

Karlan, Dean, Ryan Knight, and Christopher Udry. 2012. "Hoping to win, expected to lose: Theory and lessons on micro enterprise development." National Bureau of Economic Research.

Karlan, D., J. Morduch, and S. Mullainathan. 2010. "Take up: Why microfinance take-up rates are low and why it matters." Financial Access Initiative.

McKenzie, D. 2012. "Beyond baseline and follow-up: the case for more T in experiments." *Journal of Development Economics*.

McKenzie, David, and Christopher Woodruff. 2008. "Experimental evidence on returns to capital and access to finance in Mexico." *The World Bank Economic Review*, 22(3): 457–482.

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

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

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

Park, A. 2006. "Risk and household grain management in developing countries." *The Economic Journal*, 116(514): 1088–1115.

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

Saha, A., and J. Stroud. 1994. "A household model of on-farm storage under price risk." *American Journal of Agricultural Economics*, 76(3): 522–534.

Stephens, E.C., and C.B. Barrett. 2011. "Incomplete credit markets and commodity marketing behaviour." *Journal of Agricultural Economics*, 62(1): 1–24.

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

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

World Bank. 2006. "Malawi Poverty and Vulnerability Assessment: Investing in our Future."

Appendix A

Chapter 1

A.1 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.[1] Moreover, while I do see small price differences for purchases on credit, this relationship disappears when controlling for other features of the transaction.[2]

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.

[1] 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).

[2] 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).

Table A.1: **P**roduct 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

A.2 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 A.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.[3]

[3]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.

Figure A.1: **P**rice 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.*

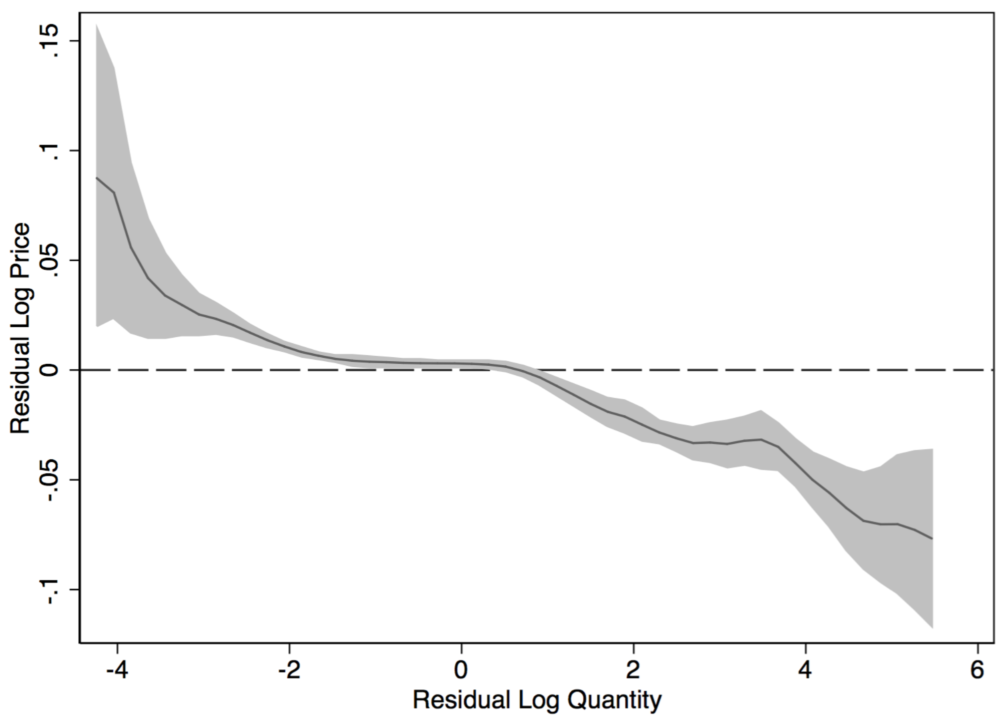

A.3 Sample Selection

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.[4]

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.[5] 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.[6]. In a large number of these markets, it was found that the information provided by the Director of Trade was inaccurate.[7] Markets that were deemed ineligible upon visit were then replaced with market from their same stratum.[8] 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.

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

[5]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).

[6]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.

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

[8]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.

A.4 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.[9] 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 A.2 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.[10] 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.[11] 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 1.4). Clear communication of the duration of the subsidy, which resulted in well-understood start and end dates, likely also contributed to this clean effect.[12]

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.[13] The magnitude of the long-run effects are a little more than half the size of the main effect, a

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

[10] 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.

[11] 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).

[12] 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.

[13] 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.

percent breakdown that is remarkably similar to that seen in Table B.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.

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.[14] 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 2.4 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.

[14]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.

Table A.2: **L**ong-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

A.5 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.[15] 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 A.3 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 1.10, 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%.[16] 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

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

[16] 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.

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.[17] Comparing these effect to the expected IV effects in Panel 1 of Table 2.4 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 A.3: **E**volution 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

[17]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.

Appendix B

Chapter 2

B.1 How do Entrants Compare to Incumbents?

This appendix provides greater detail on how incumbents compare to entrants in their market. In Table B.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:

$$Y_{idw} = \alpha + \beta E_{idw} + \lambda_{dw} + \epsilon_{idw} \tag{B.1}$$

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 B.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 B.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 B.1: **C**omparison 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 B.2: **E**ffect of Entry on Incumbents-Only. *Column 1 presents the main IV specification from Table 2.3, 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

B.2 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).

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 B.3 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 B.3: **D**uration 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 2.6:

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

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

Figure B.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 B.1: **E**ffect 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 B.3.*

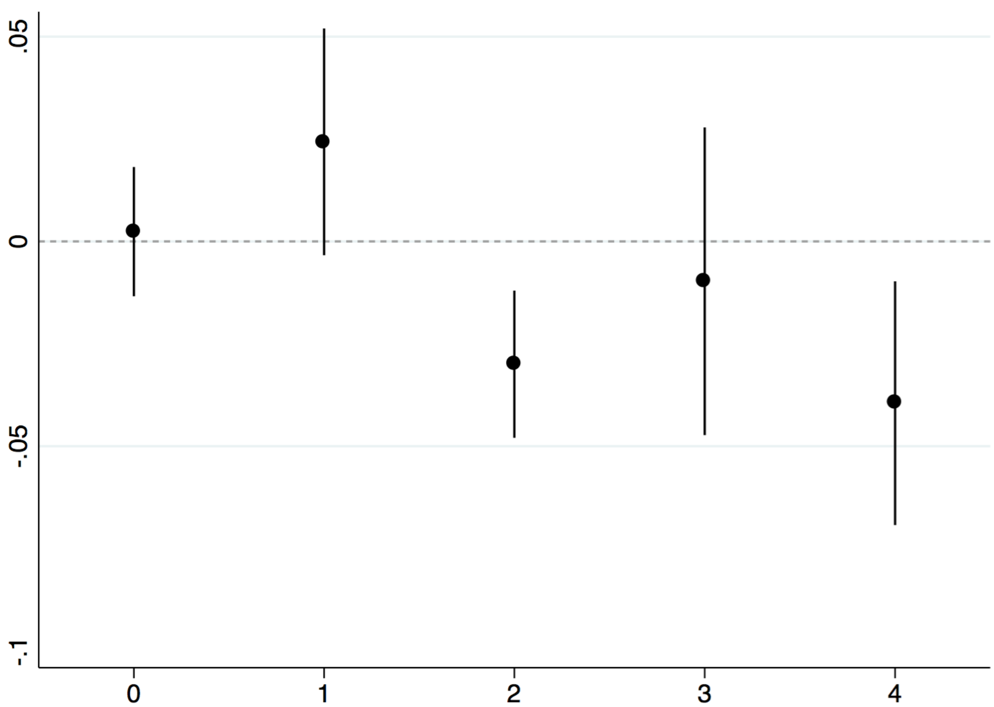

Intensity

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 B.4: **I**ntensity 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 2.6:

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

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

Figure B.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 B.2: **E**ffect 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 B.4.*

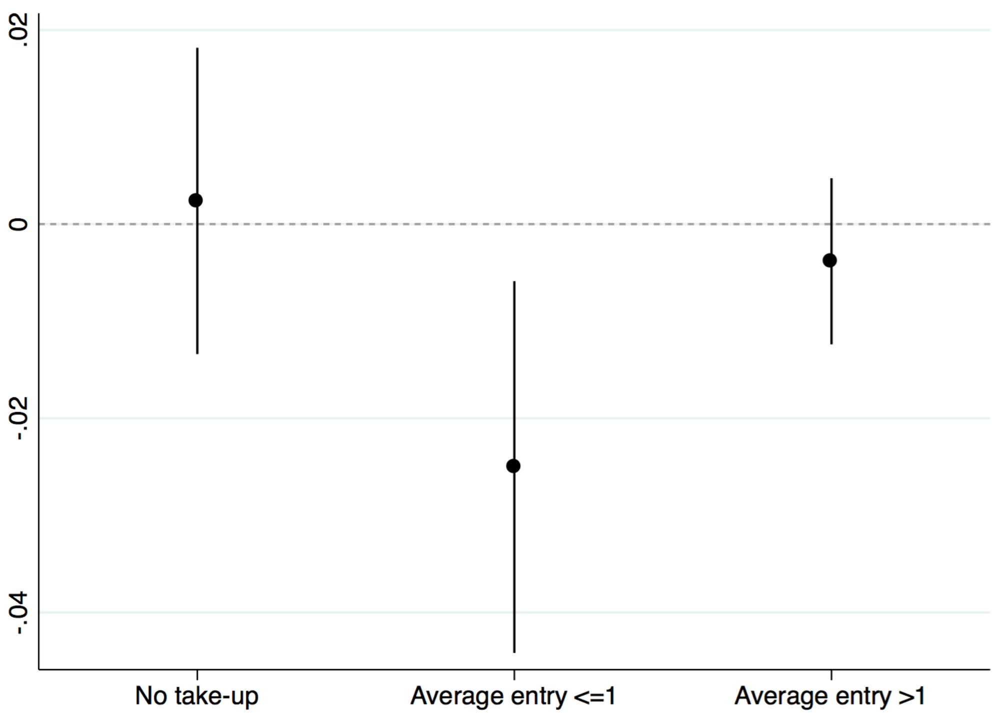

Appendix C

Chapter 3

C.1 Pilot Results

Figure C.1: **P**ilot data on maize inventories and marketing decisions over time, using data from two earlier pilot studies conducted with One Acre Fund in 2010/11 with 225 farmers (top row) and 2011/12 with 700 different farmers (bottom row). *Left panels*: inventories (measured in 90kg bags) as a function of weeks past harvest. The dotted line is the sample median, the solid line the mean (with 95% CI in grey). *Right panels*: average net sales position across farmers over the same period, with quantities shown for 2010/11 (quantity sold minus purchased) and values shown for 2011/12 (value of all sales minus purchases).

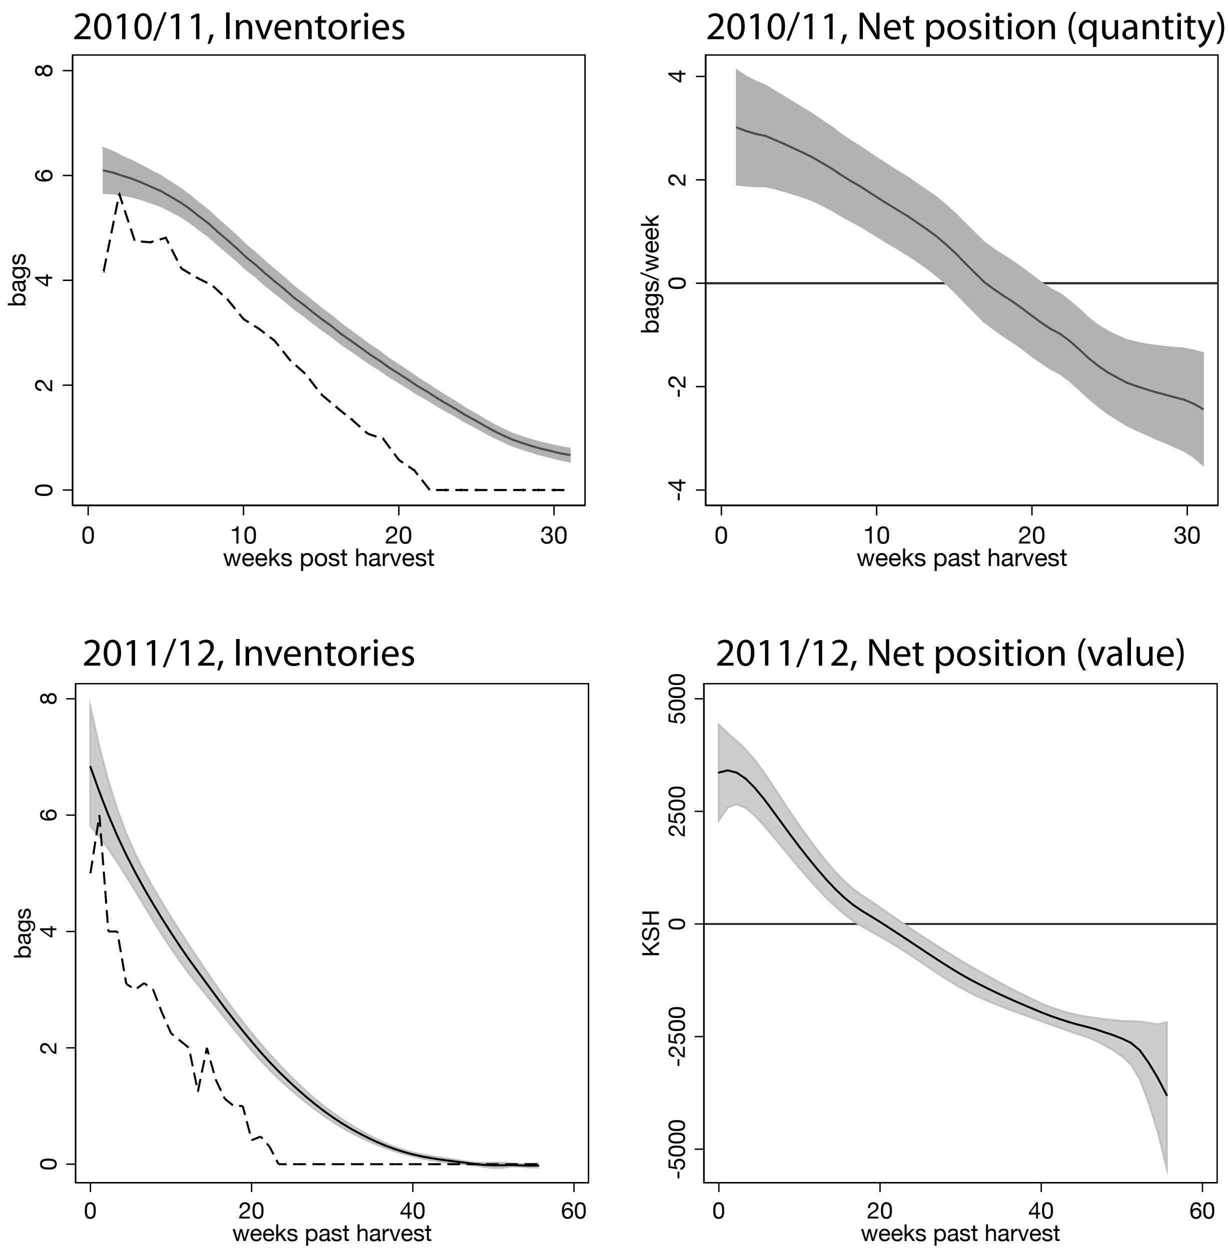

C.2 Effects of Loan Timing

In Year 1, the loan was (randomly) offered at two different times: one in October, immediately following harvest (T1) and the other in January, immediately before school fees are due (T2). Splitting apart the two loan treatment arms in Year 1, results provide some evidence that the timing of the loan affects the returns to capital in this setting. As shown in Figure C.2 and Table C.1, point estimates suggest that those offered the October loan held more in inventories, reaped more in net revenues, and had higher overall consumption. Overall effects on net revenues are about twice as high as pooled estimates, and are now significant at the 5% level (Column 5 of Table C.1), and we can reject that treatment effects are equal for T1 and T2 ($p = 0.04$). Figure C.3 shows non-parametric estimates of differences in net revenues over time among the different treatment groups. Seasonal differences are again strong, and particularly strong for T1 versus control.

Why might the October loan have been more effective than the January loan? Note that while we are estimating the intent-to-treat (ITT) and thus that differences in point estimates could in principle be driven by differences in take-up, these latter differences are probably not large enough to explain the differential effects. For instance, "naive" average treatment effect estimates that rescale the ITT coefficients by the take-up rates (70% versus 60%) still suggest substantial differences in effects between T1 and T2. A more likely explanation is that the January loan came too late to be as useful: farmers in the T2 group were forced to liquidate some of their inventories before the arrival of the loan, and thus had less to sell in the months when prices rose. This would explain why inventories began lower, and why T2 farmers appear to be selling more during the immediate post-harvest months than T1 farmers. Nevertheless, they sell less than control farmers during this period and store more, likely because qualifying for the January loan meant carrying sufficient inventory until that point.

Figure C.2: **T**reatment effects by loan timing, assuming no spillovers. Plots shows how average inventories, net revenues, and log per capita consumption evolve over the study period for farmers assigned to T1 (blue line), T2 (red line), and C (black dashed line), as estimated with fan regressions.

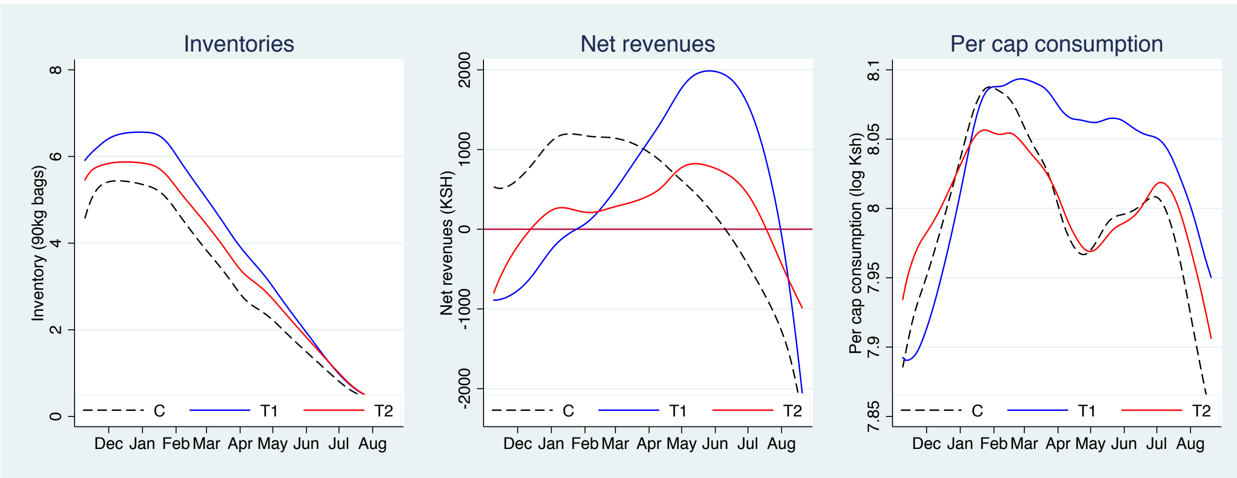

Figure C.3: **R**evenue treatment effects by loan timing, assuming no spillovers. Plots show the difference in net revenues over time for T1 versus C (left), T2 versus C (center), and T1 versus T2 (right), with bootstrapped 95% confidence intervals shown in grey.

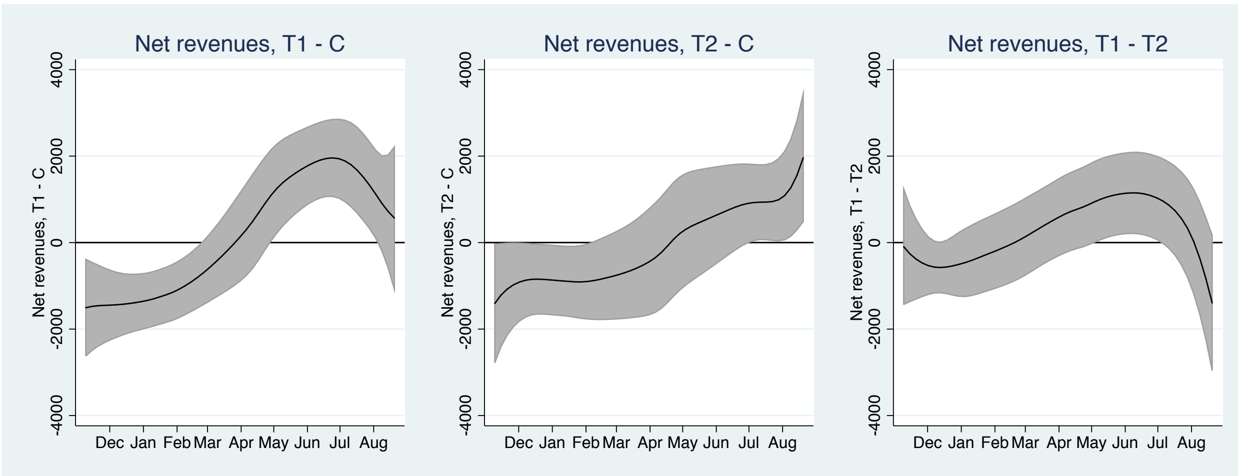

Table C.1: **Y**ear 1 Results by Loan Timing. Regressions include round fixed effects and strata fixed effects, with errors clustered at the group level.

	(1) Inventories	(2) Prices	(3) Revenues	(4) Consumtpion
T1	0.77***	-47.81**	541.95**	0.04
	(0.13)	(23.20)	(248.78)	(0.03)
T2	0.46***	2.47	36.03	0.01
	(0.13)	(22.47)	(248.15)	(0.03)
Observations	3816	1914	3776	3596
Mean of Dep Variable	3.03	2936.14	501.64	8.02
SD of Dep Variable	3.73	425.20	6217.09	0.66
R squared	0.49	0.30	0.13	0.21
T1 = T2 (pval)	0.02	0.04	0.04	0.19

C.3 Admin Price Data Mapping

In this section, we use an alternative method of matching matching markets to sublocations (and therefore to randomized treatment intensities) as a robustness check. Instead of using GPS coordinates for farmers and markets, we ask OAF field staff which markets are in their sublocation. Results in Year 1 are robust to matching method, though in Year 2 first stage results are weaker and reduced formed results are less precise using this alternate method.

Table C.2: **F**irst stage. Observed treatment intensity as a function of assigned treatment intensity. "Observed" treatment intensity is the number of treated farmers in our sample divided by the number of farmers the OAF roster for that site (as counted in Y1 for both years; this is done to avoid confounding OAF growth with changing treatment densities, since what we really care about is the percentage of total population in the area, which likely did not change much across years).

	Y1		Y2		Pooled	
	(1) Admin	(2) Nearest	(3) Admin	(4) Nearest	(5) Admin	(6) Nearest
x	0.38***	0.38***	0.13*	0.18*	0.25***	0.27***
	(0.06)	(0.05)	(0.07)	(0.08)	(0.07)	(0.07)
Observations	13	14	12	13	25	27
Mean of Dep Var	0.29	0.29	0.14	0.15	0.21	0.22
R squared	0.80	0.82	0.26	0.30	0.34	0.38

Figure C.4: **Y**1 Market prices for maize as a function of local treatment intensity (admin). The left panel shows the average sales price in markets in high-intensity areas (solid line) versus in low-intensity areas (dashed line) over the study period. The right panel shows the average difference in log price between high- and low-intensity areas over time, with the bootstrapped 95% confidence interval shown in grey. Markets to treatment intensity using administrative data on sublocation.

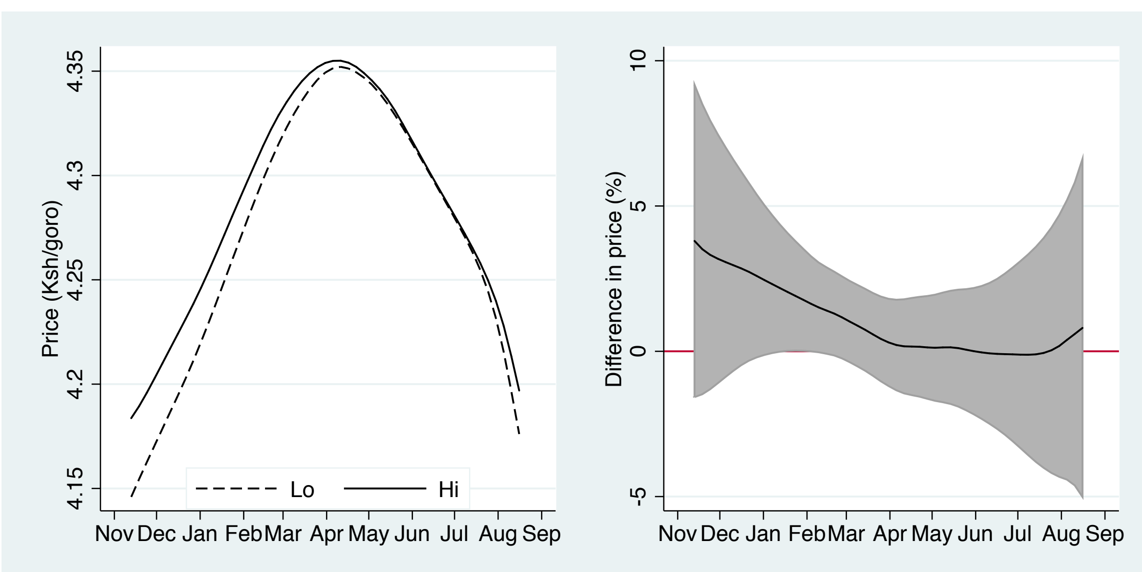

Figure C.5: **Y**2 Market prices for maize as a function of local treatment intensity (admin). The left panel shows the average sales price in markets in high-intensity areas (solid line) versus in low-intensity areas (dashed line) over the study period. The right panel shows the average difference in log price between high- and low-intensity areas over time, with the bootstrapped 95% confidence interval shown in grey. Markets to treatment intensity using administrative data on sublocation.

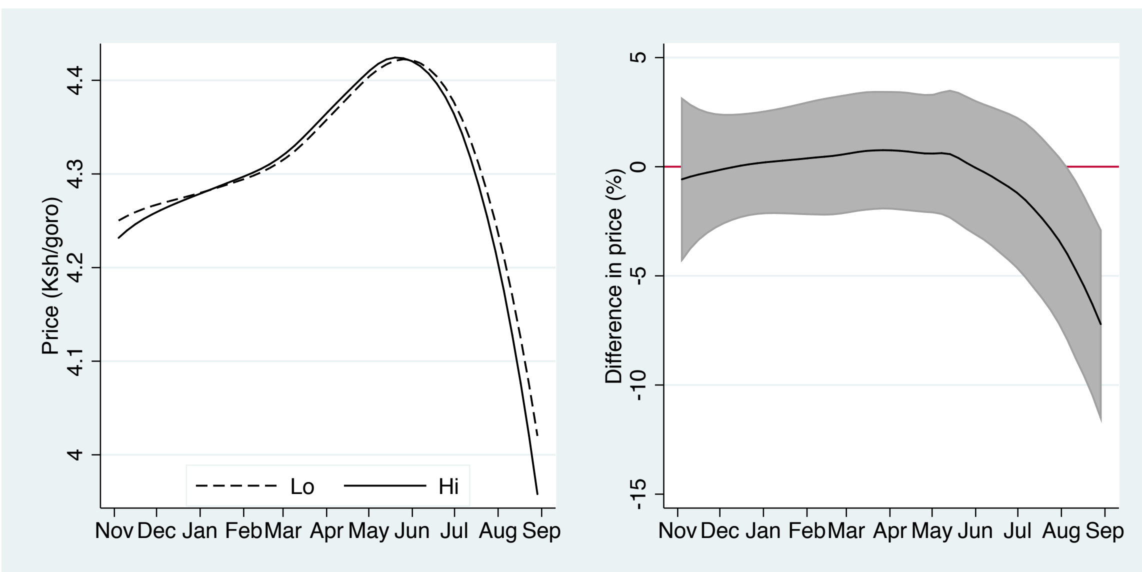

Figure C.6: **P**ooled Market prices for maize as a function of local treatment intensity (admin). The left panel shows the average sales price in markets in high-intensity areas (solid line) versus in low-intensity areas (dashed line) over the study period. The right panel shows the average difference in log price between high- and low-intensity areas over time, with the bootstrapped 95% confidence interval shown in grey. Markets to treatment intensity using administrative data on sublocation.

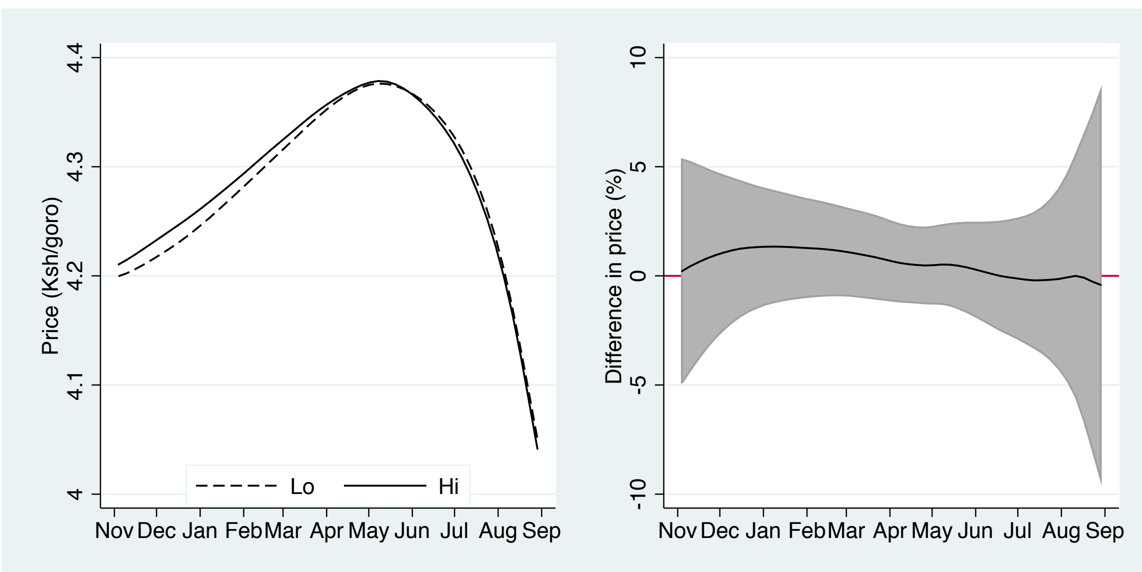

Table C.3: **P**ooled market prices for maize as a function of local treatment intensity (admin).

	Y1		Y2		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
Hi	2.64*	2.62*	0.08	0.05	1.50	1.50
	(1.25)	(1.31)	(1.16)	(1.13)	(0.92)	(0.99)
Time	0.73***	0.73***	0.84***	0.84***	0.81***	0.81***
	(0.22)	(0.22)	(0.18)	(0.18)	(0.17)	(0.17)
Hi Intensity * Time	-0.37	-0.37	-0.08	-0.08	-0.24	-0.24
	(0.27)	(0.27)	(0.26)	(0.26)	(0.22)	(0.22)
Observations	491	491	454	454	945	945
Mean of Dep Var	62.00	62.00	62.00	62.00	62.00	62.00
R squared	0.07	0.08	0.06	0.06	0.06	0.06
Controls	No	Yes	No	Yes	No	Yes

Data are for November through August in Y1 and Y2. "Hi intensity" is a dummy for a sublocation randomly assigned a high number of treatment groups and "Time" is month number (beginning in November at 0 in each year). Standard errors are clustered at the sublocation level. Markets are matched to sublocations using administrative data. Controls are the distance to the nearest road.

Table C.4: **P**ooled market prices for maize as a function of local treatment intensity. IV regression (admin).

	Y1		Y2		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
Observed Intensity	15.03**	15.03*	1.03	1.04	12.59*	12.82
	(7.34)	(7.69)	(19.86)	(18.58)	(7.60)	(8.46)
Observed Intensity * Time	-2.29	-2.28	-1.60	-1.60	-2.24	-2.23
	(1.77)	(1.77)	(3.72)	(3.77)	(1.97)	(1.97)
Time	0.95**	0.95**	0.90*	0.90*	0.98***	0.98***
	(0.42)	(0.42)	(0.47)	(0.48)	(0.37)	(0.37)
Observations	444	444	417	417	861	861
Mean of Dep Var	67.01	67.01	67.01	67.01	67.01	67.01
R squared	0.06	0.06	0.04	0.04	0.04	0.05
Controls	No	Yes	No	Yes	No	Yes

Data are for November through August in Y1 and Y2. Assigned hi low treatment intensity is used as instruments for observed intensities, where the observed intensity is the number of treated farmers divided by the number of OAF farmers in that sublocation. "Time" is month number (beginning in November at 0 in each year). Standard errors are clustered at the sublocation level. Markets are matched to sublocations using administrative data. Controls are the distance to the nearest road.

C.4 Secondary Outcomes

Table C.5: **P**ooled Non-Farm Profit Non-farm Profit is the household's profit from non-farm activities in the last month (Ksh).

	Y1		Y2		Pool	
	(1)	(2)	(3)	(4)	(5)	(6)
	Overall	By Intensity	Overall	By Intensity	Overall	By Intensity
Treat	197.30	-150.81	-127.45	-309.72	-35.28	-264.58
	(170.57)	(333.30)	(164.75)	(299.21)	(127.06)	(231.41)
Hi		-145.48		-28.99		-55.22
		(308.27)		(256.84)		(208.13)
Treat * Hi		489.84		256.78		323.31
		(385.60)		(357.42)		(275.02)
Observations	1305	1305	2938	2938	4243	4243
Mean DV	984.02	1056.54	1359.52	1337.37	1270.51	1269.33
R squared	0.00	0.00	0.00	0.00	0.00	0.00

Table C.6: **P**ooled Non-Farm Hours Hours Non-Farm is the number of hours worked by the household in a non-farm businesses run by the household in the last 7 days.

	Y1		Y2		Pool	
	(1)	(2)	(3)	(4)	(5)	(6)
	Overall	By Intensity	Overall	By Intensity	Overall	By Intensity
Treat	1.40	0.73	0.77	-0.67	0.96	-0.25
	(1.59)	(2.70)	(1.23)	(2.08)	(0.99)	(1.66)
Hi		2.40		1.14		1.41
		(2.76)		(1.69)		(1.44)
Treat * Hi		0.84		2.04		1.69
		(3.32)		(2.56)		(2.02)
Observations	1305	1305	2942	2942	4247	4247
Mean DV	11.90	10.27	13.60	12.49	13.20	11.95
R squared	0.00	0.00	0.00	0.01	0.00	0.01

Table C.7: **S**alaried Employment. Hours Salary is the total number of hours worked by household members in a salaried position.

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	0.47	0.86	0.18	-2.07	0.30	-0.96
	(1.42)	(2.43)	(1.16)	(2.18)	(0.90)	(1.64)
Hi		0.17		-1.71		-1.16
		(2.52)		(1.87)		(1.51)
Treat * Hi		-0.56		3.29		1.82
		(2.99)		(2.55)		(1.94)
Observations	1295	1295	2012	2012	3307	3307
Mean DV	11.16	10.70	6.74	7.33	8.12	8.35
R squared	0.00	0.00	0.01	0.01	0.01	0.01

Table C.8: **A**verage Wage Avg Wage is the average monthly wage for those household members who are salaried.

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	2293.22	-908.20	-333.47	1822.23	1296.43	-743.50
	(1720.62)	(3043.72)	(1620.91)	(4063.57)	(1243.91)	(2251.83)
Hi		-1843.78		-1092.62		-1476.21
		(2710.81)		(2678.38)		(1939.28)
Treat * Hi		4556.76		-2495.62		2933.25
		(3640.89)		(4689.26)		(2759.66)
Observations	284	284	135	135	419	419
Mean DV	11486.64	12087.50	5232.03	5682.00	8984.80	9278.07
R squared	0.02	0.02	0.02	0.04	0.10	0.10

Table C.9: **P**ooled School Fees Paid. School Fees Paid are the expenditure on school fees over the past month (Ksh).

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	150.82	31.71	213.27	-329.82	191.55	-94.21
	(118.32)	(227.21)	(377.33)	(693.94)	(186.63)	(335.90)
Hi		-272.68		-662.03		-485.39
		(207.59)		(573.79)		(312.27)
Treat * Hi		178.21		773.26		414.02
		(264.46)		(830.63)		(396.15)
Observations	3867	3867	2905	2905	6772	6772
Mean DV	1217.27	1369.71	3851.29	4077.54	2560.84	2740.01
R squared	0.05	0.05	0.03	0.03	0.09	0.09

Table C.10: **P**ooled Happiness Index. Happy is an index for the following question: "Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)?"

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	0.07**	0.04	0.01	0.03	0.04*	0.03
	(0.03)	(0.06)	(0.03)	(0.05)	(0.02)	(0.04)
Hi		-0.03		-0.02		-0.02
		(0.06)		(0.04)		(0.04)
Treat * Hi		0.04		-0.03		0.01
		(0.07)		(0.06)		(0.05)
Observations	3870	3870	2969	2969	6839	6839
Mean DV	2.57	2.58	2.68	2.68	2.63	2.63
R squared	0.01	0.01	0.00	0.00	0.01	0.01

C.5 Long-Run Follow-up (LRFU) Survey Results

The Long-Run Follow-Up (LRFU) survey was run Nov-Dec 2015. Results presented in this appendix show the limited effects of the loan on long-run outcomes.

Table C.11: **L**RFU 2014-2015 Outcomes: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Harvests and amount eaten are in 90kg bag units.

	2014 Harvest			Amt. Eaten		
	Y1	Y2	Both	Y1	Y2	Both
Y1	-0.51		-1.06	-0.23		-0.46
	(0.45)		(0.86)	(0.23)		(0.45)
Y2		0.51	0.45		-0.06	-0.20
		(0.44)	(0.96)		(0.24)	(0.51)
Y1 * Y2			0.64			0.24
			(1.15)			(0.63)
Observations	973	935	556	969	933	553
R squared	0.00	0.00	0.01	0.00	0.00	0.01
Mean DV Control	9.02	8.86	9.61	5.51	5.60	6.14

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.12: **L**RFU 2014-2015 Outcomes: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Net sales are in 90kg bag units. Revenues are in Ksh.

	Net Sales			Revenues		
	Y1	Y2	Both	Y1	Y2	Both
Y1	0.31		-0.01	350.50		-763.60
	(0.35)		(0.59)	(950.10)		(1854.40)
Y2		0.29	0.29		1286.62	1330.40
		(0.35)	(0.61)		(1094.42)	(1777.33)
Y1 * Y2			0.21			1126.71
			(0.80)			(2510.70)
Observations	979	937	557	979	938	558
R squared	0.00	0.00	0.00	0.00	0.00	0.01
Mean DV Control	-0.10	0.35	0.46	397.23	1052.01	1422.30

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.13: **L**RFU 2014-2015 Sales: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) sales. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Amounts are in 90 kg bag units and values are in Ksh.

	Tot Amt Purch			Tot Val Purch		
	Y1	Y2	Both	Y1	Y2	Both
Y1	0.10		0.01	557.96		252.96
	(0.23)		(0.47)	(645.31)		(1363.64)
Y2		0.17	-0.12		338.96	-236.18
		(0.22)	(0.55)		(670.48)	(1534.45)
Y1 * Y2			0.29			773.24
			(0.67)			(1893.17)
Observations	979	935	555	979	936	556
R squared	0.00	0.00	0.00	0.00	0.00	0.00
Mean DV Control	2.01	2.13	2.26	5646.07	6342.74	6387.60

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.14: **L**RFU 2014-2015 Purchases: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) sales. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Amounts are in 90 kg bag units and values are in Ksh.

	Tot Amt Sold			Tot Val Sold		
	Y1	Y2	Both	Y1	Y2	Both
Y1	0.08		0.20	298.39		407.96
	(0.15)		(0.25)	(452.26)		(726.89)
Y2		-0.23	-0.33		-811.94	-1274.11
		(0.17)	(0.28)		(531.18)	(792.22)
Y1 * Y2			0.13			829.11
			(0.35)			(1010.21)
Observations	978	938	557	978	938	557
R squared	0.01	0.01	0.01	0.01	0.01	0.01
Mean DV Control	1.90	1.86	1.72	5560.79	5590.23	5220.76

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.15: **L**RFU 2014-2015 Sales Harvest Season: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) sales. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Amounts are in 90 kg bag units and values are in Ksh.

	Harv Amt Sold			Harv Val Sold		
	Y1	Y2	Both	Y1	Y2	Both
Y1	0.03		0.25	77.46		530.06
	(0.09)		(0.16)	(243.35)		(481.72)
Y2		0.18**	0.22		334.68	600.49
		(0.08)	(0.21)		(221.93)	(603.64)
Y1 * Y2			-0.22			-572.62
			(0.24)			(707.79)
Observations	980	937	555	980	935	556
R squared	0.00	0.00	0.01	0.00	0.00	0.00
Mean DV Control	0.52	0.46	0.36	1346.28	1267.63	1079.90

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.16: **L**RFU 2014-2015 Sales Lean Season: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) sales. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Amounts are in 90 kg bag units and values are in Ksh.

	Lean Amt Sold			Lean Val Sold		
	Y1	Y2	Both	Y1	Y2	Both
Y1	0.22		0.16	679.47		392.38
	(0.20)		(0.41)	(574.93)		(1155.90)
Y2		0.04	0.06		303.79	115.41
		(0.20)	(0.49)		(568.03)	(1307.93)
Y1 * Y2			0.05			513.65
			(0.60)			(1676.81)
Observations	981	937	557	981	935	557
R squared	0.00	0.00	0.00	0.00	0.00	0.00
Mean DV Control	1.34	1.53	1.49	3974.15	4383.35	4354.60

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.17: **L**RFU 2014-2015 Purchases Harvest Season: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) purchases. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Amounts are in 90 kg bag units and values are in Ksh.

	Harv Amt Purch			Harv Val Purch		
	Y1	Y2	Both	Y1	Y2	Both
Y1	-0.04		0.17	-149.68		347.10
	(0.09)		(0.15)	(233.61)		(375.77)
Y2		-0.08	-0.01		-298.29	-146.51
		(0.08)	(0.17)		(215.31)	(406.71)
Y1 * Y2			-0.19			-370.52
			(0.20)			(494.05)
Observations	977	941	557	977	940	557
R squared	0.01	0.00	0.02	0.01	0.00	0.02
Mean DV Control	0.58	0.52	0.44	1484.23	1317.58	1144.34

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.18: **L**RFU 2014-2015 Purchases Lean Season: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) purchases. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Amounts are in 90 kg bag units and values are in Ksh.

	Lean Amt Purch			Lean Val Purch		
	Y1	Y2	Both	Y1	Y2	Both
Y1	0.10		-0.03	370.11		-294.98
	(0.12)		(0.20)	(356.84)		(628.83)
Y2		-0.09	-0.31		-279.60	-1092.92
		(0.13)	(0.21)		(416.36)	(668.77)
Y1 * Y2			0.34			1432.54
			(0.27)			(869.14)
Observations	982	939	559	979	938	558
R squared	0.00	0.01	0.01	0.00	0.01	0.01
Mean DV Control	1.29	1.25	1.27	3922.78	3926.80	4040.25

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.19: **L**RFU 2015 Harvest and Input Use: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on 2015 LR harvest. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Harvests are in 90kg bag units. Results are for maize plots only.

	2015 Harvest		
	Y1	Y2	Both
Y1	-0.22		-1.53*
	(0.56)		(0.92)
Y2		0.92	-0.42
		(0.59)	(0.94)
Y1 * Y2			2.39*
			(1.27)
Observations	987	946	561
R squared	0.00	0.00	0.02
Mean DV Control	9.78	9.97	10.95

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.20: **L**RFU 2015 Harvest and Input Use: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on 2015 LR input usage. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Non-labor input exp are the amount spent in Ksh on all fertilizers, hybrid seeds, DAP, CAN, and other physical inputs excluding labor. Labor person-days record the number of person-days of labor applied. Results are for maize plots only.

	Non-Labor Input Exp			Labor Person-Days		
	Y1	Y2	Both	Y1	Y2	Both
Y1	18.46		315.04	-4.76		-13.76
	(213.39)		(393.59)	(5.98)		(9.85)
Y2		122.23	-153.46		-9.66	-16.38
		(194.98)	(404.36)		(7.04)	(13.00)
Y1 * Y2			402.65			14.63
			(526.04)			(15.84)
Observations	978	940	559	979	940	560
R squared	0.01	0.00	0.01	0.01	0.00	0.06
Mean DV Control	2620.61	2271.07	2001.67	126.15	131.48	142.58

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.21: **L**RFU 2015 Food Consumption and Expenditure: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on food consumption and expendiure. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Maize Eaten in the past week in 2kg "goros." Food expenditure is the value of maize purchases, own production consumed, and gifts given to others over the past 30 days.

	Maize Eaten			Food Exp		
	Y1	Y2	Both	Y1	Y2	Both
Y1	-0.11		0.43	40.82		-124.26
	(0.19)		(0.38)	(247.76)		(492.87)
Y2		-0.26	-0.13		99.58	-97.26
		(0.22)	(0.41)		(251.35)	(556.87)
Y1 * Y2			-0.47			254.32
			(0.54)			(658.28)
Observations	976	937	554	977	939	557
R squared	0.00	0.00	0.01	0.02	0.00	0.02
Mean DV Control	5.68	5.74	5.51	6840.11	6786.12	6928.43

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.22: **L**RFU 2015 Consumption and Happiness: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on consumption and happiness. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* HH consumption is the total household consumption (logged) over the past 30 days. Happy is an index for the following question: "Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)?"

	HH Cons			Happy		
	Y1	Y2	Both	Y1	Y2	Both
Y1	-0.03		-0.00	0.10**		0.05
	(0.05)		(0.10)	(0.05)		(0.08)
Y2		0.04	0.08		0.01	0.00
		(0.05)	(0.11)		(0.04)	(0.10)
Y1 * Y2			-0.09			-0.03
			(0.13)			(0.12)
Observations	976	939	556	985	945	560
R squared	0.01	0.00	0.01	0.01	0.00	0.01
Mean DV Control	9.50	9.47	9.49	2.40	2.47	2.48

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.23: **L**RFU 2015 Education: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment education and non-farm profit. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Education attendance is the proportion of days the children in the household attended school in the last 5 days. Educational expenditure is the total household expenditure on children's education (in Ksh) over the past 12 months.

	Edu Exp			Edu Attend		
	Y1	Y2	Both	Y1	Y2	Both
Y1	-3654.14		-6576.46	0.00		0.02
	(3854.68)		(6998.49)	(0.01)		(0.02)
Y2		-1168.61	-4367.33		-0.01	0.02
		(2917.71)	(8041.06)		(0.01)	(0.02)
Y1 * Y2			2391.45			-0.04
			(9231.27)			(0.03)
Observations	979	936	556	927	876	528
R squared	0.00	0.00	0.01	0.00	0.00	0.01
Mean DV Control	38371.63	37452.55	43373.16	0.94	0.95	0.93

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.24: **L**RFU 2015 Non-Farm Business Hours Worked and Profits: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on non-farm business. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Hours Non-Farm is the number of hours worked by the household in a non-farm businesses run by the household in the last 7 days. Non-farm profit is the household's profit from non-farm activities in the last month (Ksh).

	Hours Non-Farm			Non-Farm Profit		
	Y1	Y2	Both	Y1	Y2	Both
Y1	0.94		1.41	-186.29		48.03
	(1.75)		(2.71)	(285.72)		(528.13)
Y2		0.22	0.63		-244.86	-47.72
		(1.87)	(3.43)		(315.71)	(607.26)
Y1 * Y2			4.05			-47.91
			(4.25)			(744.40)
Observations	979	937	556	975	933	552
R squared	0.01	0.00	0.02	0.00	0.00	0.01
Mean DV Control	15.97	14.87	13.32	2138.25	2019.84	1966.83

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.25: **L**RFU 2015 Salaried Employment Hours and Wages: Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on salaried employment. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* Hours Salary is the total number of hours worked by household members in a salaried position. Avg Wage is the average monthly wage for those household members who are salaried.

	Hours Salary			Avg Wage		
	Y1	Y2	Both	Y1	Y2	Both
Y1	-2.28		1.47	1892.96		884.26
	(1.77)		(3.57)	(1697.63)		(3231.62)
Y2		-0.98	-1.74		3651.39**	528.77
		(1.98)	(4.49)		(1700.71)	(3525.65)
Y1 * Y2			-4.57			3027.24
			(5.19)			(4752.24)
Observations	982	939	559	292	274	155
R squared	0.00	0.00	0.01	0.00	0.02	0.02
Mean DV Control	15.03	14.30	15.50	13014.88	12646.63	12714.71

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table C.26: **L**RFU Market prices for maize as a function of local treatment intensity.

	(1) Admin	(2) Admin	(3) Nearest	(4) Nearest
Hi	-0.26	-0.30	-0.78	-0.78
	(1.76)	(1.72)	(1.37)	(1.31)
Time	1.72***	1.72***	1.61***	1.61***
	(0.17)	(0.17)	(0.21)	(0.21)
Hi Intensity * Time	-0.14	-0.11	0.13	0.14
	(0.45)	(0.45)	(0.41)	(0.40)
Observations	261	261	261	261
Mean of Dep Var	62.00	62.00	61.87	61.87
R squared	0.25	0.25	0.25	0.25
Controls	No	Yes	No	Yes

Data are for November 2014 through August 2015 (the year following the cessation of Y2 treatment). Assigned treatment intensities are used as instruments for observed intensities, where the observed intensity is the number of treated farmers divided by the number of OAF farmers in that sublocation. "Hi intensity" is a dummy for a sublocation randomly assigned a high number of treatment groups and "Time" is month number. Standard errors are clustered at the sublocation level. Columns 1 and 2 match markets to sublocations using administrative data, columns 3 and 4 using location data on farmers and markets. Controls are the distance to the nearest road.

C.6 Price Effects Balance and Robustness

To further check robustness of the price results in Year 1, we start by dropping sublocations one-by-one and re-estimating prices differences. As shown in the left panel of Figure C.7, differential trends over time in the two areas do not appear to be driven by particular sublocations. Second, building on other experimental work with small numbers of randomization units (Bloom et al., 2013; Cohen and Dupas, 2010), we generate 1000 placebo treatment assignments and compare the estimated price effects under the "true" (original) treatment assignment to estimated effects under each of the placebo assignments.[1] Results are shown in the two right hand panels of Figure C.7. The center panel shows price differences under the actual treatment assignment in black, and the placebo treatment assignments in grey. "Exact" p-values on the test that the price difference is zero are then calculated by summing up, at each point in the support, the number of placebo treatment estimates that exceed the actual treatment estimate and dividing by the total number of placebo treatments (1000 in this case); these are shown in the right-hand panel of the figure. Calculated this way, prices differences are significant at conventional levels for the first 3-4 months post harvest, roughly consistent with the results shown in Figure 3.5.

Figure C.7: **R**obustness of price effect estimates. *Left panel*: difference in prices between high and low-density markets over time for the full sample (black line) and for the sample with each sublocation dropped in turn (grey lines). *Center panel*: price effects under the "true" treatment assignment (black line) and 1000 placebo treatment assignments (grey lines). *Right panel*: randomization-inference based p-values on the test that the price difference is zero, as derived from the center panel.

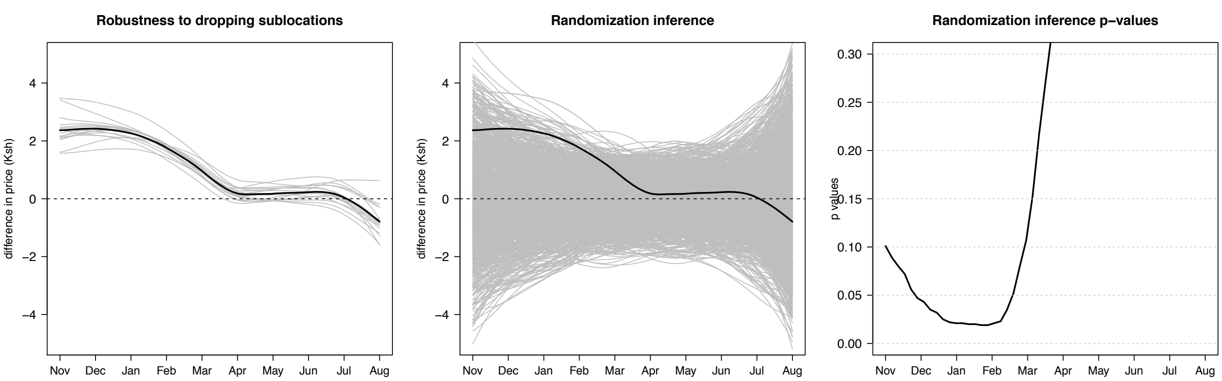

[1] With 17 sublocations, 9 of which are "treated" with a high number of treatment farmers, we have 17 choose 9 possible treatment assignments (24,310). We compute treatment effects for a subset of these possible placebo assignments.

Table C.27: **B**alance among baseline covariates, high versus low treatment intensity areas. The first two columns give the means in the low or high treatment intensity areas, the 3rd column the total number of observations across the two groups, and the last two columns the differences in means normalized by the standard deviation in the low intensity areas, with the corresponding p-value on the test of equality.

	Lo	Hi	Obs	Lo - Hi	
				sd	*p-val*
Male	0.32	0.31	1,589	0.02	0.72
Number of adults	3.11	3.07	1,510	0.02	0.74
Kids in school	3.15	2.98	1,589	0.09	0.11
Finished primary	0.71	0.75	1,490	-0.08	0.13
Finished secondary	0.27	0.25	1,490	0.04	0.51
Total cropland (acres)	2.60	2.35	1,512	0.08	0.15
Number of rooms in hhold	3.31	3.08	1,511	0.08	0.10
Total school fees (1000 Ksh)	29.23	27.88	1,589	0.04	0.51
Average monthly cons (Ksh)	15,586.03	14,943.57	1,437	0.05	0.38
Avg monthly cons./cap (log Ksh)	7.98	7.97	1,434	0.02	0.77
Total cash savings (KSH)	5,776.38	6,516.09	1,572	-0.04	0.56
Total cash savings (trim)	5,112.65	4,947.51	1,572	0.01	0.82
Has bank savings acct	0.42	0.42	1,589	-0.01	0.91
Taken bank loan	0.07	0.09	1,589	-0.06	0.30
Taken informal loan	0.25	0.24	1,589	0.02	0.72
Liquid wealth	87,076.12	98,542.58	1,491	-0.12	0.06
Off-farm wages (Ksh)	3,965.65	3,829.80	1,589	0.01	0.84
Business profit (Ksh)	1,859.63	2,201.34	1,589	-0.04	0.53
Avg %Δ price Sep-Jun	121.58	138.18	1,504	-0.21	0.00
Expect 2011 LR harvest (bags)	10.52	8.70	1,511	0.08	0.03
Net revenue 2011	-2,175.44	-4,200.36	1,428	0.03	0.45
Net seller 2011	0.34	0.30	1,428	0.08	0.16
Autarkic 2011	0.06	0.07	1,589	-0.04	0.53
% maize lost 2011	0.01	0.01	1,428	0.00	0.95
2012 LR harvest (bags)	11.57	10.94	1,484	0.07	0.19
Calculated interest correctly	0.68	0.74	1,580	-0.12	0.03
Digit span recall	4.49	4.60	1,504	-0.10	0.08
Maize giver	0.25	0.27	1,589	-0.05	0.37
delta	0.14	0.13	1,512	0.07	0.28

See Table 3.1 and the text for additional details on the variables.

C.7 Effect of Tags

In this section, we test whether loan treatment effects are actually being driven by the tags. To test this, in Year 1 we offered a small additional treatment arm in which farmers received only the tags. Table C.28 presents the comparison between the overall treatment effect from the full intervention in Year 1 to that just of the tags.

We see in Table C.28 that point estimates are larger across the board for the pooled and T1 groups than for the tags-alone group (which are never significant). That said, estimates are somewhat noisy, due to the small sample size of the tags-alone group, and only for inventories can we formally reject that the effect of the loan was driven by the tags.

Table C.28: **E**ffects of tags. Regressions include round fixed effects and strata fixed effects, with errors clustered at the group level.

	(1) Inventories	(2) Inventories	(3) Revenues
Treat Y1	0.52***	276.29	0.00
	(0.16)	(291.42)	(0.03)
Tags	0.06	71.00	-0.01
	(0.23)	(411.42)	(0.05)
Observations	4273	4229	4223
R squared	0.34	0.01	0.00
pooled-tags p-val	0.06	0.63	0.89

C.8 Attrition

This section explores the differential attrition from the Year 1 to the Year 2 sample.

Table C.29: **S**elective attrition Did those who stayed in the sample have different outcomes variables or treatment effects? Year 1 outcome variables regressed on dummy for whether treated in Y1, dummy for whether stayed in the sample in Y2, and interaction term. Sample is all Y1 subjects. Treatment effects at the individual level, all rounds. Regressions include round fixed effects and strata fixed effects, with errors clustered at the group level.

	(1)	(2)	(3)	(4)	(5)
	Invent	Purchase price	Sales prices	Rev	Log HH Cons
Treat Y1	0.60***	9.88	-19.52	315.01	-0.01
	(0.14)	(23.90)	(27.55)	(302.74)	(0.04)
Stayed Y2	0.51***	78.91**	-44.41	380.62	0.01
	(0.19)	(31.38)	(39.10)	(338.42)	(0.05)
Treat Y1 * Stayed Y2	-0.08	-100.28**	43.30	-158.30	0.05
	(0.23)	(38.84)	(45.41)	(408.97)	(0.06)
Observations	3816	1914	1425	3776	3792
Mean of Dep Variable	2.67	2982.02	2827.58	334.41	8.00
R squared	0.50	0.30	0.47	0.13	0.03
Controls	Yes	Yes	Yes	Yes	Yes

Table C.30: **A**ttrition and Sample Selection. "Attrit" is an indicator for having exited the sample between Year 1 (2012-13) and Year 2 (2013-14). "Stay" is an indicator for being in the Year 1 and Year 2 samples

Baseline characteristic	Attrit	Stay	Obs	Attrit - Stay *sd*	*p-val*
Treatment 2012	0.56	0.66	1,589	-0.20	0.00
Male	0.28	0.25	1,816	0.07	0.13
Number of adults	3.01	3.12	1,737	-0.05	0.30
Kids in school	2.89	3.23	1,816	-0.17	0.00
Finished primary	0.73	0.77	1,716	-0.08	0.10
Finished secondary	0.25	0.25	1,716	-0.01	0.81
Total cropland (acres)	2.26	2.50	1,737	-0.08	0.12
Number of rooms in hhold	2.94	3.34	1,738	-0.16	0.00
Total school fees (1000 Ksh)	25.93	30.08	1,816	-0.11	0.02
Average monthly cons (Ksh)	14,344.56	15,410.58	1,652	-0.09	0.10
Avg monthly cons./cap (log Ksh)	7.94	7.96	1,649	-0.04	0.49
Total cash savings (KSH)	5,355.05	6,966.35	1,797	-0.09	0.13
Total cash savings (trim)	4,675.61	4,918.86	1,797	-0.02	0.70
Has bank savings acct	0.38	0.46	1,816	-0.15	0.00
Taken bank loan	0.07	0.08	1,816	-0.04	0.46
Taken informal loan	0.23	0.24	1,816	-0.01	0.86
Liquid wealth	89,564.21	100,021.77	1,716	-0.10	0.05
Off-farm wages (Ksh)	3,508.17	4,103.66	1,816	-0.05	0.31
Business profit (Ksh)	2,069.13	2,159.55	1,816	-0.01	0.86
Avg %Δ price Sep-Jun	130.30	141.63	1,728	-0.15	0.00
Expect 2011 LR harvest (bags)	8.13	9.55	1,732	-0.09	0.05
Net revenue 2011	-4,983.94	-4,156.75	1,633	-0.02	0.72
Net seller 2011	0.26	0.35	1,633	-0.19	0.00
Autarkic 2011	0.06	0.07	1,816	-0.03	0.53
% maize lost 2011	0.01	0.01	1,609	0.00	0.98
2012 LR harvest (bags)	9.26	11.94	1,708	-0.31	0.00
Calculated interest correctly	0.72	0.72	1,806	-0.01	0.91
Digit span recall	4.61	4.50	1,731	0.09	0.06
Maize giver	0.26	0.26	1,816	0.00	0.98
Delta	0.14	0.13	1,738	0.08	0.09