Price Information, Inter-Village Networks, and "Bargaining Spillovers": Experimental Evidence from Ghana

Nicole Hildebrandt, Yaw Nyarko, Giorgia Romagnoli, Emilia Soldani *

30 July 2014

Preliminary draft: please do not cite without authors' permission.

Abstract

This paper presents results from a randomized experiment designed to evaluate the impact of providing commodity price information to farmers in rural Ghana. We find substantial positive effects of the intervention: prices for treatment group farmers increased by about 11%, an effect that is sustained over the longer term (2 years). Moreover, we find that control group farmers with strong network ties to the treatment group also realized large, indirect benefits comparable in size to the direct effect on the treated. The indirect effects on control group prices are not explained by information sharing between treatment and control farmers. Instead, we present and provide evidence for a novel mechanism — which we term "bargaining spillovers" — through which these indirect benefits occurred.

JEL Codes: D82, O13, Q11, Q12, Q13.
Keywords: Price information, Agriculture, Bargaining, ICTs, Networks, Externalities.

*New York University and CTED (Center for Technology and Economic Development). We are extremely grateful to NYU-Abu Dhabi Institute, as well as anonymous donors, for generous financial funding, without which this project could not have taken place. We are also very grateful to Isaac Boateng and our team of field interviewers for logistical support. We thank audiences at the American University of Sharjah, NYU-Abu Dhabi, the 2013 African Econometrics Society Summer Meeting, and the 2014 North American Econometrics Society Summer Meeting for many useful comments and suggestions. The authors' email contacts are: nth211@nyu.edu; yaw.nyarko@nyu.edu; gr763@nyu.edu; and es1945@nyu.edu.

1 Introduction

The rapid increase in mobile phone coverage and ownership in developing countries is making it easier to provide farmers with accurate, (near) real-time information on prices to help them make optimal marketing decisions. Can such market information help farmers get higher prices for their production? And, what are the indirect impacts of information provision on traders, on farmers that do not have access to price information, and on market outcomes as a whole? These are important questions to answer, given the growing interest in ICT-related informational interventions by policymakers, foundations, and governments around the world.

This paper reports results from a two-year randomized evaluation of an SMS-based market information system (MIS) in Ghana. Our study involved 1,000 smallholder commercial farmers in the northern part of the Volta region, who we followed for two years between 2011 and 2013. The intervention consisted of providing treatment group farmers with a subscription to a price alert service that sent weekly text messages with local and urban market prices for their main commercial crops. We find that price information had significant benefits on treatment farmers' prices, and that the effect is sustained over the longer term (2 years). Beyond documenting these direct effects, our paper makes two key contributions to the existing literature.[1] First, we use information we collected on inter-village marketing and communication networks to show that, over time, control group farmers with strong network ties to the treatment group realized large indirect benefits from the intervention. Second, we develop and provide evidence for a novel mechanism (which we call "bargaining spillovers") through which these indirect benefits occurred.

We focus on evaluating the impact of the price alerts on yam prices, the most important crop in our study area. We start by estimating treatment effects under the standard RCT assumption that the intervention had no impact on the prices received by the control group (i.e there are no spillovers). We find a positive and statistically significant short-term effect of the price alerts on yam prices, but no effect in the longer run. There are two possible explanations behind these results: (1) our "no spillovers" assumptions is correct, and there is no long-run benefit of the MIS;

[1] Camacho and Conover (2011), Fafchamps and Minten (2012), Mitra, Mookherjee, Torero and Visaria (2013), Svensson and Yanagizawa (2009), and Nakasone (2013) study the impacts of similar price alert systems in Colombia, Maharastra, West Bengal, Uganda, and Peru, respectively. A paper by Goyal (2010) studies a related intervention involving information kiosks in district markets in Andhra Pradesh. There is also a literature looking more broadly at the impact of mobile phone coverage on agricultural outcomes in the developing world; see Jensen (2007), Aker (2008), Aker and Fafchamps (2010), and Muto and Yamano (2009).

or (2) our assumption is violated because control farmers start to indirectly benefit from the service. In the first case, an estimated treatment effect that is not significantly different from zero truly reflects zero impact, while in the second case, indirect benefits for control group farmers creates downward bias in the estimated treatment effect. This happens because control group prices are converging upward to treatment group prices, and thus no longer represent the counterfactual of interest (what treatment prices would have been absent the intervention).

Using data we collected on inter-village marketing and communication networks, we provide evidence of substantial indirect benefits on control group farmers' yam prices due to the intervention. We show that, over time, control farmers with stronger network ties to the treatment group increasingly benefited from higher yam prices relative to control farmers with weaker network ties to the treatment group. Next, we attempt to determine the mechanism through which these indirect effects are generated. Interestingly, our data provide little empirical support for the most obvious potential mechanism: information sharing between treatment and control group farmers. To the extent that information spillovers are present, they are weak and do not appear to be the driving force behind our results.

As a next step, we develop a model of farmer-trader bargaining with asymmetric information that suggests a different channel through which indirect benefits can arise. In our model of "bargaining spillovers," traders are unable to perfectly distinguish between informed and uninformed farmers. Instead, they form a prior on the likelihood that a given farmer is informed. Traders know that informed farmers will reject low offers, so if the trader's prior about the farmer is high enough, she offers him a higher price, regardless of whether or not he is actually informed. While most of the model's predictions are indistinguishable from an information spillovers story, there is one prediction that cannot be rationalized by information spillovers alone, and we find empirical support for this prediction in our data. Taken together, our results support the hypothesis that bargaining spillovers are the main mechanism causing the indirect benefits for control group farmers.

In the final section of the paper, we present estimates of the corrected, or de-biased, treatment effects (i.e. accounting for indirect effects on control group prices). In the second year of the study, the direct impact on treatment farmers was 17.62 GHS per 100 tubers of yam, an 11.4% increase in prices.[2] On top of this, we find large indirect impacts on control group farmer prices, amounting

[2]Percentage increases in this paragraph are calculated based on mean prices for the control group, exclusive of

to about 17.18 GHS per 100 tubers on average (an 11.1% increase). For the typical yam farmer selling 1,200 tubers in a year, these effects translate into an additional 200 GHS (US$135) in annual revenue. Given that farmers tend to operate on low margins, the impact is likely to be considerably larger in terms of a percentage increase in farm profits. For example, assuming a profit margin of 50% and no change in costs, an 11% increase in prices is equivalent to a 22% increase in profits. From an ROI perspective, in comparison to the costs of the service and training, the direct return from purchasing an annual subscription exceeds 200%. The ROI is even higher if the indirect benefits are also accounted for.

In addition to contributing new evidence on the impact of MIS on agricultural outcomes, our paper is related to the growing body of experimental research that evaluates the indirect effects of interventions in realms where externalities are likely to occur, such as disease control (Miguel and Kremer, 2004), labor markets (Crépon, Duflo, Gurgand, Rathelot, and Zamora, 2013), and elections (Asunka, Brierley, Golden, Kramon, Ofosu, 2014; Giné and Mansuri, 2012). In the realm of agriculture, few experimental studies have ventured beyond looking for evidence of informational spillovers among farmers in the same village.[3] Our paper finds large indirect effects across villages, through a mechanism other than information sharing. Our results suggest that indirect impacts of interventions in agricultural markets can be substantial, and are therefore extremely important to take into consideration.

The remainder of this paper is structured as follows. Section 2 provides an overview of agricultural marketing in Ghana and outlines our experimental design. Section 3 describes the data we collected during the study and presents some descriptive statistics. Section 4 presents our estimates of the treatment effect under the assumption of zero spillovers. Section 5 presents evidence of indirect benefits for control group farmers, and Section 6 presents the model of bargaining spillovers and the supporting empirical evidence. Section 7 presents estimates of de-biased treatment effects. Section 8 concludes.

spillovers. All figures in this paragraph are denominated in real, August 2011 Ghana Cedis.

[3]One paper that has ventured beyond looking for evidence of information spillovers is Svensson and Yanagizawa-Drott (2012), which looks at the partial and general equilibrium effects of a national MIS in Uganda. In their theoretical model and subsequent empirical analysis, uninformed farmers lose out in general equilibrium, due to reductions in urban market prices (and corresponding reductions in farm-gate prices) brought about by changes in the marketing behaviors of informed farmers. The distributional impacts implied by their results are quite different from our findings: in our case, due to spillovers, uninformed farmers not hurt—and in fact are likely benefiting—from the MIS. The differences between this paper and ours are likely attributable to the differing scales of each MIS as well as differences in the marketing environment being studied.

2 Background and Experimental Design

2.1 Agricultural marketing in Ghana

As in other parts of sub-Saharan Africa, farmers in Ghana heavily rely on traders (middlemen) to market their production. This is largely a consequence of the country's poor transportation infrastructure and urban population concentration in the south, factors that make it difficult for farmers to directly access final consumer markets. Traders are individuals (often women) who travel around the country purchasing agricultural output from farmers, and then transport this output to urban markets to sell.[4] Transactions between farmers and traders usually take place at the farm gate, in the local community, or in the local market. They are conducted in an informal manner (formal contractual relationships are rare)[5] and involve some amount of bargaining between the parties. The degree to which bargaining takes place varies by crop, as described further below.

Because traders travel extensively, they tend to have detailed knowledge on market prices and trends, significantly more so than farmers.[6] Farmers often complain that they are being cheated by traders, and cite examples of traders telling them urban prices are low (which farmers are unable to verify) in order to buy at a low price. Given the information asymmetry that exists between farmers and traders, one potentially viable way to help farmers secure higher prices is by providing them with better price information. The success of this type of intervention depends critically on the industrial organization of the trader side of the market: in a perfectly competitive trader market, an information asymmetry should not affect farmers' outcomes; while in a monopolistic trader market, solving the information asymmetry will not have any effect on farmers' outcomes. In our study area, traders are best described as an oligopoly. Barriers to entry are high, since trading requires access to capital and a network of farmers with which to transact. On average, farmers report having 2 to 4 different traders with which to sell in a given agricultural season, many (2 to 3) of whom are repeat buyers with whom the farmer has a long-term relationship. Given this environment, price information has the potential to have a positive impact on farmers' bargaining

[4]There are also small-scale, local traders, who aggregate crops to re-sell to the large traders. Our focus is on the large traders, since they are the dominant actors in the market and the individuals that farmers identify as cheating them in negotiations.

[5]See, e.g. Quartey et al, 2012.

[6]This is borne out by the fact that farmers in our study cite traders as a primary source of information on market prices.

outcomes with traders.

Details on the way in which farmers and traders transact vary considerably by type of crop. Some key differences are illustrated in Table I. One crop that stands out across several dimensions is yam, the crop that is the focus of our study. Yam is the only crop that is sold in urban markets by a non-negligible proportion of farmers, although even for yam, direct sales to urban markets are relatively uncommon.[7] Yam is also the only crop for which bargaining is a universal feature of crop marketing. For products such as maize and gari (a form of processed cassava), prices are fairly homogenous among sellers in the local market, and farmers often report paying the prevailing "market price" for their production. This is not the case for yam, a crop that farmers told us has no reference "market price." Instead, the farmer's ability to successfully negotiate with the trader is a crucial determinant of the final price. Another disparity we observed during our field work is that most yam trading takes place the day before the actual market day, in a separate area of the marketplace.[8] These distinctive features of yam marketing are likely contributing explanations to our finding that the MIS had a significant effect for yams, but no effect for other crops.

2.2 Study location and sample selection

We conducted our experiment in the northern part of the Volta region, an area that lies in central-eastern Ghana, approximately 300km from Accra.[9] Within the study area, we sampled 100 communities located in four contingous districts: Krachi East, Krachi West, Nkwanta North, and Nkwanta South.[10] From each community, we sampled 10 farmers to be included in the study among those who market at least some portion of his or her crop (i.e. we excluded subsistence-only farmers). Nine farmers declined to be part of the study, leaving our final sample at 991 farmers.

[7] In the baseline, 24% of farmers reported selling yam in urban markets, and 15% reported urban markets as their main place of sale.

[8] We also observed that large-scale traders were significantly more active in the purchase of yam than they were of other crops.

[9] We chose this area for two reasons. First, the area is "virgin territory" in the sense that the MIS we study was not previously present in this area, and there are few NGOs operating there. Second, the area is fairly self-contained geographically: the Togo border lies to the east, and the Volta Lake lies to the west.

[10] Ghana consists of 10 administrative regions, which are further subdivided into districts. There are approximately 216 districts in the entire country, and 25 districts within the Volta region. Communities are identified by the central government as concentrations of people that live within a well-defined geographic area.

2.3 Randomization strategy

Our randomization strategy was developed to achieve two goals. First, we sought to minimize the indirect (spillover) effects of the intervention, in particular the risk of informational contamination (i.e. the treatment group sharing the price alerts with the control group). Second, we sought to achieve balance, meaning that we wanted the treatment group and control group to be as similar as possible prior to the introduction of the treatment.

A well known tradeoff exists between the two goals: minimizing spillovers requires that treatment and control groups be sufficiently far apart geographically, while balance requires that treatment and control groups be similar to each other, and similarity usually calls for geographical proximity (Duflo, Glennerster, and Kremer, 2007). Our solution is a two step strategy: we randomize at the "community cluster" level (to minimize spillovers) and use a stratified randomization procedure (to ensure balance).

In RCTs where information spillovers are a concern, randomization is often carried out at the community level to minimize the risk of information contamination. This design implicity assumes that information transmission is likely to occur among individuals *within* the same community, but is unlikely to occur *across* communities. We were reluctant to follow this approach, since our preliminary field work suggested that some communities in our sample were quite well connected to each other. Instead, we opted for a design that groups highly-connected communities together into what we call a "community cluster," and then randomizes at the community cluster level. Our process of forming community clusters is briefly summarized here, and described in greater detail in Appendix A.

Our baseline survey gathered information on inter-village marketing and communication networks. We used this information to construct three indices for each village pair j and k:

1. *Market overlap index*: measuring the degree to which farmers in villages j and k sell in the same markets

2. *Marketing communications index*: measuring the degree of communication about agricultural marketing between farmers in villages j and k

3. *Geographic proximity index*: measuring the geographic distance between two villages j and k

We used principal components analysis to create a single "connectedness index" out of the three indices listed above. The connectedness index provides a scalar measure of connectedness, c_{jk}, for each village pair j and k in the study. Higher values denote more connected village pairs, and lower values denote less connected village pairs. To form community clusters, we selected a cut-off value for c_{jk}, above which village pairs were put into the same cluster, and below which village pairs were kept in separate clusters. In order to preserve balance and power, we chose a fairly low cut-off value, which resulted in moving from 100 communities to 90 community clusters.

In addition to informing our randomization, the connectedness index allows us to investigate the indirect effects of the intervention ex-post. As discussed in Section 5, we use this index to construct a measure of each community's "connectedness" to the treatment group (what we call "C2T"), to investigate these indirect effects.

With only 90 community clusters, we were concerned that a simple randomization could result in an imbalance between the treatment and control groups. To mitigate this risk, we opted for a stratified randomization procedure. We stratified the community clusters on district (Nkwanta North, Nkwanta South, Krachi East, Krachi West), and most commonly-grown crop (yam, or not yam). This resulted in 8 strata (four districts by two crop categories). Within each strata, we randomly assigned half of the community clusters to the treatment group, and half to the control group. The procedure resulted in 45 clusters (49 villages) in the treatment group and 45 clusters (51 villages) in the control group.

2.4 Details about the treatment

Farmers in the treatment group were trained and given a free subscription to an MIS operated by a privately-held technology company called Esoko. The MIS provides weekly price alerts to subscribers via SMS (text message).[11] We registered farmers to receive price alerts for their two main commercial crops, for four local markets in the study region and four of the main urban markets in the country.[12] Enrolled farmers started receiving weekly price alerts in late-October

[11] Esoko relies on a network of "market enumerators" to collect these market prices. Esoko trains enumerators to ensure that prices are collected in a consistent manner across markets, and holds twice-yearly refresher trainings to keep enumerators sharp. In addition, the company quality reviews all prices before they are sent out and occasionally employs "mystery shoppers" to validate the information sent in by enumerators. Esoko operates its MIS in 16 countries across the African continent.

[12] The four urban markets are Accra-Agbogbloshie, Accra-Ashaiman, Tema, and Koforidua. The four local markets are Nkwanta, Kpassa, Borae, and Dambai. Prior to the start of our experiment, Esoko did not monitor prices at

2011.[13] Since most markets in the country are weekly, this should in theory provide farmers with the most up-to-date price information available.

Farmers in the control group were not provided with trainings or a subscription to the price alert service. However, they were surveyed with the same frequency as farmers in the treatment group. Our surveys are described in more detail in the next section.

3 Data

Our analysis relies on two types of data: monthly data on sale and annual data on background characteristics. The monthly data, carried out from August 2011 through June 2013, gather detailed information about every sales transaction conducted by the farmer for his two main commercial crops. The information we collect in these surveys includes: quantity and variety sold, total revenue, price per unit, place of sale, and type of buyer. Annual data consist of three surveys: (i) a baseline survey conducted in July-August 2011 (prior to the start of the intervention); (ii) a midline survey conducted in July-August 2012 (about nine months after the start of the intervention); and (iii) an endline survey conducted June-August 2013 (about 1.5 years after the start of the intervention). These surveys cover a wide range of topics, including marketing environment, price information and aggregate data on sales for farmers' two main crops.[14]

The richness of our data allows us to provide new empirical evidence on the impact of MIS along two dimensions. First, using the monthly data we are able to compare short- and longer-run effects and look at the *dynamics* of the treatment over time. Second, the detailed information in the annual data allows us to test competing hypothesis for the mechanisms which originated such dynamics and ultimately inform a toy model of bargaining spillovers. It turns out that our results are of interest along both of these dimensions, in that they help to reveal the indirect effects of the MIS on control group farmers.

these local markets, due to the fact that it had virtually no MIS subscribers in this area. As part of the study, we commissioned Esoko to begin gathering these market prices.

[13]Price alerts were in English, one of Ghana's official languages. Prices were sent in local unit measures, e.g. 100 tubers of yam, 1 long bag of maize.

[14]Most farmers in Ghana grow a variety of crops for consumption and sale, rather than focus exclusively on a single crop. This is also true in our sample.

3.1 Descriptive Statistics and Balance

Table II reports baseline summary statistics for the full sample and separately by treatment status, as well as tests for balance between the treatment and control groups. Overall, the variables are well balanced between treatment farmers and control farmers. With only two exceptions, the p-values on the difference in means between the treatment and control groups are not statistically significant, meaning we cannot reject that the means are the same for the two groups.

In the full sample, farmers are 41 years old on average, are predominantly male, and rely on farming as the main source of household income. The sample is not highly educated: while 42% have completed junior high school, nearly 50% have no formal education. Median income earned from the farmer's two main commercial crops amounted to GHS 1,400 (US$898) in the agricultural season ending in June 2011.[15]

The main crops grown by farmers in the sample are yam, cassava, maize, and groundnut. Yam is by far the most commonly grown crop, with over 60% of farmers reporting it as one of their two main commercial crops. Farmers' knowledge of urban market prices is very low: only about 30% of farmers believe that they are well informed about urban market prices at the time of the baseline survey. Farmers are more informed about local market prices, which is consistent with the finding that the majority sell in local markets.

4 Impact on prices under the assumption of no spillovers

To measure the impact of the price alerts, we start by estimating the treatment effect under the Stable Unit Treatment Value Assumption (SUTVA) typically invoked in RCT evaluations. This assumption says that the potential outcomes of an individual are unrelated to the treatment status of any other individual. Under SUTVA, we assume that there are no spillovers from the treatment that end up affecting the prices of control group farmers. With this assumption, we can estimate the causal effect of the price alerts by using the monthly sales data to estimate the following regression:

$$p_{ijt} = \lambda + \kappa T_j + X'_{ijt}\psi + \alpha_k + \alpha_t + e_{ijt} \tag{1}$$

[15] Dollar figures are calculated using the average GHS-USD exchange rate for 2011 from Oanda.com.

where p_{ijt} is the producer price outcome for farmer i living in community cluster j selling in month t, T_j is a treatment status indicator, α_k are strata fixed effects, α_t are period fixed effects, and X_{ijt} is a set of additional covariates. The coefficient κ estimates the effect of the price alerts on an Intent-to-Treat (ITT) basis. It is an unbiased estimate of the treatment effect so long as SUTVA is not violated (as it would be in the presence of spillovers). We include strata fixed effects and period fixed effects in our regressions, so that the treatment effect is identified from within-period, within-strata variation between treatment and control groups.

We estimate (1) separately for Year 1 (November 2011-June 2012) and Year 2 (July 2012-June 2013). The results using Year 1 data provide an estimate of the short-run treatment effect, and the Year 2 data provide an estimate of the longer-run treatment effect. We also combine all the data (including three months of pre-treatment data, from August 2011-October 2011) to estimate the following pooled regression:

$$p_{ijt} = \sum_{s=0}^{2} \{\lambda_s Y_s + \kappa_s (T_j * Y_s)\} + X'_{ijt}\psi + \alpha_k + \alpha_t + e_{ijt} \tag{2}$$

where Y_s is an indicator for period s, $s \in \{0, 1, 2\}$ (pre-treatment, Year 1, and Year 2). In this regression, the λ_s measure the average control group price in period s, and the κ_s measure the average treatment effect in period s. As in (1), the κ_s are unbiased estimates of the treatment effect on an ITT basis, so long as SUTVA holds.

We estimate both (1) and (2) separately for each crop, due to the fact that there are important differences in the marketing environment across crops, and therefore there is likely to be treatment effect heterogeneity. Here, we focus on the results for yam, which is the most important crop in our study. It is also the only crop for which we find any evidence of a treatment effect. The fact that we only find an effect for yam is most likely explained by two factors. First, there is evidence that markets in Ghana are fairly well integrated for grains such as maize and rice. Where markets are well integrated, there may be less of a need for market price information in negotiations with buyers. Second, as we show in Appendix B, increases in yam prices are mainly attributable to improvements in farmers' bargaining position with traders, rather than to changes in where crops are sold or the timing of sales. However, while bargaining is a central element of yam marketing, it is much less important for other crops.

4.1 Results

Table III presents results from the estimation of (1) and (2) using the monthly sales data. The dependent variable in these regressions is yam prices per 100 tubers, denominated in real August 2011 Ghana Cedis (GHS). The first two columns present results using the data from Year 1, the second two columns present results using data from Year 2, and the final two columns present results using the pooled data. For each cut of data, we present results from two different specifications: one that only controls for strata fixed effects, period fixed effects, and yam type; and a second that also includes additional covariates (gender, asset index, and the community's distance to the closest district market).[16]

Consistently across our specifications, we find a positive and statistically significant treatment effect in Year 1 (short-term effect). According to column (6), the estimated treatment effect in the first year was about 8.69 Ghana Cedis per 100 tubers, significant at the 5% level. This amount represents about 6.5% of the control group's mean price in Year 1. In contrast, the longer-term treatment effect is small in magnitude and never significantly different from zero. Again looking at column (6), the estimated treatment effect in Year 2 is 0.433 (not significant), or about 0.2% of the control group's mean price.

We also examined the impact of the treatment on quantities of yam sold, and found no significant impact at any point in time. Thus, our estimated treatment effects on prices can also be thought of as treatment effects on farmers' revenues from farming. Given the initial positive treatment effect, the next step is to understand the mechanisms through which this occurred. Our analysis of the mechanisms is presented Appendix B, where we show that the main mechanism is through changes in farmers' bargaining with traders.

Taken together, our results suggest that, under the assumption of zero spillovers, there was initially a positive treatment effect of the alerts, but it disappeared in the long run. We explore this further by using non-parametric methods (fan regressions) to look at the evolution of yam

[16]In mid-2012, we discovered that the surveyors in the Nkwanta North district were falsifying some of the data in the monthly surveys. Rather than go back and have the work redone in retrospect, we decided to simply discard the suspect data. Thus, the monthly data relied on in this paper does not contain information for Nkwanta North from August 2011 through June 2012. Given our stratified sampling approach, the omission of this data should not distort any of results, although it does reduce sample size which may lead to greater imprecision in some of our estimates. Results using the annual data (where we did not have to drop any data) are comparable, and are available upon request.

prices for the treatment and control groups over time. The results of this analysis are presented in Figure I. The top figure plots yam prices for the treatment and control groups, as estimated using fan regressions after controlling for strata, type of yam, and additional covariates.[17] The bottom figure plots the difference between the treatment and control groups, with the bootstrapped 95% confidence interval shown in grey. The figure demonstrates that, immediately following the introduction of the price alerts, there is a large estimated treatment effect: over 20 GHS per 100 tubers, more than twice the estimated short-term effect in Table III. The treatment effect steadily declines over time so that five to six months after the start of the intervention, it is no longer significantly different from zero. The pattern displayed in the bottom panel of the figure suggests that the disappearance of the treatment effect is not a consequence of seasonality in the yam marketing season. Why, then, did the effect disappear?

5 Explaining the treatment effect dynamics

The results presented in Section 4 present an interesting puzzle, which we now attempt to better understand. Crucially, our interpretation of the estimated treatment effects hinges on whether or not SUTVA has been violated. If SUTVA holds, then our estimated treatment effects are unbiased and reflect the true long-term effect of the MIS on farmers' prices. A possible explanation behind this result is a "fade out" story: over time, farmers stopped paying attention to the alerts, which caused their prices to revert to what they would have been absent the intervention. If "fade out" occurred, then, at endline, we should see fewer treatment farmers reporting using the alerts in their marketing. However, as shown in Table IV, we see the opposite trend, casting doubt on a fade out story.[18] Another possibility is that traders stopped transacting with informed farmers. However, this explanation is inconsistent with anecdotal information we gathered in interviews with traders, who report rarely changing their trading route due to the high costs associated with entering new areas. Moreover, in our annual surveys, informed farmers did not report experiencing a reduction in trading partners at midline or endline.

The explanation that is most consistent with our data is that the intervention indirectly bene-

[17] The additional covariates are the same ones used in Table III.

[18] Farmers may be over-reporting use of Esoko simply to please us, but it is not clear why such over-reporting would be higher at endline than at midline, so it is difficult to see how over-reporting could explain the increase in the reported use of Esoko over time.

fited some control group farmers, which caused control group prices to converge upward to treatment group prices. This caused a violation of SUTVA, leading to a downward bias in the estimated treatment effects presented in Section 4. Next, we present the empirical evidence of positive spillovers on control group farmers' prices. We then explore the possible mechanisms behind these spillover effects. Interestingly, our data suggests that the most obvious candidate — information sharing between treatment and control farmers — is not the primary mechanism driving the indirect benefits for control group farmers. Instead, our data is consistent with what we call a "bargaining spillovers" story. A simple model of bargaining spillovers is presented in Section 6.1, but the basic idea is as follows. The set-up of the model assumes that farmers and traders are randomly matched with each other each period, and that traders are unable to perfectly distinguish between informed and uninformed farmers. Instead, they form a prior on the likelihood that the farmer they are matched with is informed. Traders know that informed farmers will reject low offers, so if the trader's prior about the farmer is high enough, she offers him a higher price, regardless of whether or not he is actually informed. Most of the model's predictions are indistinguishable from an information spillovers story, but there is one prediction that cannot be rationalized by information spillovers alone, and we find empirical support for this prediction in our data. Taken together, our results support the hypothesis that bargaining spillovers are the main mechanism causing the indirect benefits for control group farmers.

5.1 Empirical evidence of indirect benefits for control group farmers

In this section, we present evidence that some control group farmers indirectly benefited from the intervention. Our approach is based on the assumption that control group farmers with stronger network ties to treatment group farmers are more likely to experience indirect benefits. Here, we use the term "network ties" to refer to the degree to which farmers are connected to one another (how much they communicate and/or otherwise interact with one another).

The logic behind this assumption depends on the mechanism that is actually driving the indirect benefits for control farmers. With information spillovers, the argument is straightforward: control farmers with stronger network ties to treatment farmers are more likely to communicate about the price alerts with the treatment group. In the case of bargaining spillovers, the credibility of this assumption depends on how a farmer's network ties affect traders' perceptions on whether the farmer is informed about market prices. One possibility is that the farmer's connections to treatment group villages informs the traders' perceptions. Traders are likely to believe that farmers living in villages that are very connected to treatment group villages (through communication networks, or through selling in similar markets, or simply through distance) are also more likely to be informed about urban market prices.

Under the assumption that control farmers with stronger network ties to the treatment group are more likely to experience indirect benefits, we can test for indirect benefits by (1) constructing a measure of the strength of a farmer's network ties to treatment group farmers; and then (2) examining the relationship between network ties to the treatment group and prices. To do this, we rely on the inter-village network data collected at baseline for our randomization. We argue that the connectedness measure we created for our randomization, c_{jk}, is a good measure of the strength of network ties between farmers living in villages j and k. We use the connectedness index (c_{jk}) to construct a measure of each village's connectedness to the treatment group, which we call "C2T". For a village j, this variable is constructed as the simple average of all the c_{jk} scores for all villages $k \in T$ and $k \neq j$, rescaled to lie between zero and one:

$$\mathrm{C2T}_j = \frac{\frac{1}{N_T}\sum_{k\in T, k\neq j} c_{jk} - \min\left(\frac{1}{N_T}\sum_{k\in T, k\neq j} c_{jk}\right)}{\max\left(\frac{1}{N_T}\sum_{k\in T, k\neq j} c_{jk}\right) - \min\left(\frac{1}{N_T}\sum_{k\in T, k\neq j} c_{jk}\right)} \tag{3}$$

where N_T is the number of treatment group villages.

To estimate the relationship between C2T and prices, we calculate the C2T measure for all villages (treatment and control) and run the following regression on the monthly sales data:

$$p_{ijt} = \sum_{s=0}^{2} \{\delta_s Y_s + \alpha_s(T_j * Y_s) + \beta_s(C_j * \mathrm{C2T}_j * Y_s) + \gamma_s(T_j * \mathrm{C2T}_j * Y_s)\} + \alpha_k + \alpha_t + X'_{ijt}\psi + e_{ijt} \tag{4}$$

The outcome of interest, p_{ijt}, is the price outcome of farmer i, in community j, in month t. T_j and C_j are treatment and control group indicator variables, respectively, and $\mathrm{C2T}_j$ is as defined above. As before, the regression includes strata fixed effects (α_k) and period fixed effects (α_t), as well as other covariates. We are interested in seeing changes in the impact of C2T over time, so we interact variables with a set of time indicators Y_s, $s \in \{0, 1, 2\}$ (pre-treatment, Year 1, and Year 2).

The main coefficients of interest are the β_s, which represent the impact of C2T for control group farmers at each time period. Similarly, the γ_s coefficients represent the impact of C2T for treatment group farmers at each time period. The differential impact of C2T on control farmers relative to treatment farmers is captured by the difference $(\beta_s - \gamma_s)$. We have conjectured that positive spillovers accruing to the control group are the reason why the estimated treatment effect falls over time. In order for this to be true, we should see: (1) β_s getting larger over time and (2) $(\beta_s - \gamma_s)$ getting larger over time.

Figure II presents the key results from this regression. The top panel of the figure shows the estimated β_s and γ_s coefficients for the pre-treatment period, Year 1, and Year 2, along with 90% confidence intervals. As predicted, the β_s coefficients are increasing over time, and becomes significantly different from zero (at the 90% level) in Year 2. The γ_s coefficients seem to be moving in the opposite direction, although they are never significantly different from zero. The bottom panel shows the difference $(\beta_s - \gamma_s)$ for each time period. There is a strong upward trend in the differential impact of C2T on control group prices (relative to the treatment group), and this differential is significantly different from zero in Year 2. Taken together, the results are consistent with a spillover story being the cause of the decline in the treatment effect over time.

A potential limitation of our approach is that, besides capturing spillover effects, the C2T variable could be positively related to prices for other reasons. For example, more connected villages could have better access to markets, be more accessible to traders, and/or have better access to information from non-Esoko sources. We are not too worried about this spurious relationship, for two reasons. First, β_0 and γ_0, the estimated impacts of C2T in the pre-treatment period, are small and not significantly different from zero. Second, the impacts of C2T are very different for control and treatment group farmers over time. Because we randomized treatment status, if the relationship between C2T and prices is truly spurious, then it should be the same for both groups.

As a final piece of evidence ruling out concerns of a spurious relationship, we re-run the regression above, but this time we use measures of connectedness to *control* group villages ("C2C"). C2C is calculated in the same manner as C2T, but it considers connections to control group villages. If we are picking up a spurious relationship between connectedness and prices using the C2T measure, we should probably see a similar set of results using the C2C variable. Results for C2C are presented in Figure III, and they are quite different from the trends for C2T, particularly with regard to the differential impact on control group prices relative to treatment group prices. Using C2C, the differential is essentially flat over time, which means that C2C *cannot* explain the disappearance in the treatment effect over time.

To conclude, we find evidence that farmers in control villages with stronger network ties to treatment villages realized benefits from these network ties, which is consistent with both bargaining spillovers and information spillovers. We now focus on determining which of these mechanisms is driving the results.

5.2 Testing for information spillovers

Our annual surveys collected several measures of farmers' knowledge of market prices, which we can use to test for information sharing between the treatment group to the control group. We start by looking at farmers' (subjective) beliefs about their own knowledge of market prices. In the baseline and endline surveys, we asked farmers: "Do you feel that you are well informed about URBAN [LOCAL] market prices?"[19] Table V presents how treatment status, C2T, and their interaction are related to farmers' responses to this question. In the endline data, the treatment indicator is positive and statistically significant, indicating that treatment farmers feel more informed about market prices than control farmers.[20] The coefficients on C2T for control group farmers are not significantly different from zero, indicating that more connected control farmers do not feel more informed than less connected control farmers.

We can also look at more objective measures of price knowledge to test for information spillovers. In the baseline and midline surveys, we asked farmers to list the markets for which they regularly knew market prices. In Table VI, we look at the relationship between the number of markets listed by farmers and treatment status, C2T, and their interaction. In the midline results, treatment status has a significant, positive impact on number of markets listed. However, the estimated coefficient on C2T for the control group is not significantly different from zero, implying no statistically significant informational advantage to more connected control group farmers.

A final piece of evidence comes from the endline survey. We asked farmers to estimate current (contemporaneous) prices for yam in Accra. To see whether treatment farmers were more "correct" in their guesses, we calculated two measures of farmer estimation error: (1) absolute error =

[19]Potential answers were: (1) "no, not at all"; (2) "no, not very well"; (3) "yes, fairly well", and (4) "yes, very much."

[20]Interestingly, at baseline, treatment farmers felt *less* informed than control farmers.

abs(Accra price - farmer guess); and (2) absolute percentage error = abs(Accra price - farmer guess)/Accra price.[21] We then ran regressions of the natural logarithms of these errors on treatment status, C2T, and their interaction. As shown in Table VII, treatment farmers were more accurate in their guesses about Accra prices, a result which is significant at the 10% level.[22]. Again, the estimated coefficients on C2T for the control group are not significantly different from zero.

The results presented here provide compelling evidence that, even at the end of the study, treatment farmers were subjectively and objectively more informed about market prices than control group farmers. There is no statistically significant relationship between C2T and levels of market information for the control group, indicating that information spillovers were relatively weak (if they were present at all) in our study. Thus, they do not seem to be the driving force behind our results.

6 A model of bargaining spillovers

The data indicate that spillovers are the cause of the disappearance of the treatment effect over time. We have provided evidence that information spillovers cannot fully explain the result: statistically significant differences in farmer's levels of price knowledge remain even at the end of the survey. We have posited an alternative type of network externality, which we call *bargaining spillovers*, as explaining our results. In this section, we provide an economic model of farmer-trader bargaining with asymmetric information. The key idea captured by the model is that the provision of information has positive externalities also on farmers that do not have access to it neither directly, through the MIS, nor indirectly, through information spillovers. This is due to the fact that informed and uninformed farmers sell to the same pool of traders and traders adjust their overall bargaining strategies in the face of the new condition brought about by the MIS. A possible channel through which the bargaining externalities take place is that the traders cannot perfectly distinguish the informed farmers from the uninformed farmers. Instead they can only roughly identify areas in the network where it is plausible that the price information is received and know which villages gravitate more closely around such areas. Under this view the trader's strategy will not depend on the actual information possessed by his opponent, but rather only on

[21]We used the prices provided in the text messages to capture the "true" Accra price.

[22]See the first column of results for each measure of estimation error

some partial assessment of the information available to the farmer, which we model formally as the probability that the farmer is informed. We posit that a proxy for such probability is the index of connectedness to the treatment villages (as measured by the $C2T$ variable). This is due for example to the fact that villages that are more connected to the treatment group are more likely to receive price information, or at least it may be so believed by the traders. Moreover villages that are more connected to the treatment group are also geographically closer to them and are therefore more likely to sell to the same traders that interact with the treated farmers. The main prediction of the model is that, when the probability that a farmer is informed is high enough, the trader treats him as if he was informed with probability one. It follows that farmers with high C2T (or high probability of being informed) will receive the same price offer received by treatment farmers. The model allows us to derive one additional prediction which would be hard to reconcile with a simple theory of information spillovers. We test this prediction against the data in section 5.4.2, finding confirmation for the bargaining spillover theory.

6.1 A theory of bargaining in a network of traders and farmers with asymmetric information

We present a model of bargaining under asymmetric information, which is an adaptation of the Myerson (1984) bargaining model to a multi-period and multi-type framework. The game is in discrete time. In each period there is a continuum of identical traders with the same value v. v is *iid* over time and drawn from the uniform distribution with support $[v_L, v_H]$. The value v could represent, for example, the resale price of the commodity in the urban markets. Traders live for one period. There is a continuum of infinitely-lived farmers that have one unit of commodity for sale and discount the future by a factor β. All agents are risk neutral. In each period a farmer and a trader are matched and bargain over the price p of the commodity[23]. With probability $w \in (0, 1)$ the farmer makes a take-it-or-leave-it offer p that the trader can either accept or reject. With the remaining probability $(1 - w)$ the trader makes a take-it-or-leave-it offer to the farmer. w and

[23]The assumption that farmers are infinitely lived while traders live only for one period gives analytical tractability and, more importantly, captures a fundamental difference between farmers and traders. Farmers have a fixed supply of harvest to sell. Hence they compare the current price with the continuation value of waiting and selling in the future. Instead traders do not buy in fixed amounts. They treat each bargaining session in isolation and in each they compare the price with the resale value of the commodity in the urban market. In this sense they are modeled as short-lived. Alternatively they can be thought of as infinitely lived, but with a continuation value that does not depend on the outcome of the bargaining session currently under consideration.

$(1-w)$ capture the bargaining power of farmers and traders respectively. If the offer is accepted by the opponent, the trader's utility is $(v-p)$, while the farmer's utility is p. If the offer is rejected, the trader receives utility 0 and the farmer keeps the commodity and moves to the next period. The discounted continuation value of a farmer who does not reach an agreement and move to the next period is R. Farmers can be of two types, $i = \{I, U\}$. Informed farmers (I) know the value v, while uninformed farmers (U) only know the distribution from which v is drawn. Traders do not know whether they are matched to an informed or uninformed farmer, but know the probability d that the farmer is informed.

Optimal strategies and equilibrium

There are 3 types of agents: traders, informed farmers and uninformed farmers. Each type plays two roles: proposer and respondent. For each type we describe the optimal strategies as proposer and respondent. In Proposition 1 we characterize the unique equilibrium.

The optimal strategies of the respondents are easily determined. A trader accepts if and only if his/her value is at least as high as the price offered by the farmer, that is if and only if $v-p \geq 0$. Define as R^i, $i = \{U, I\}$, the discounted continuation value of a farmer of type i after a bargaining round ending in a rejection. A farmer of type i accepts if and only if the price offered by the trader is at least as high as R^i, that is if and only if $p \geq R^i$. We now turn to the strategy of the proposers.

Informed farmer's strategy

The informed farmer always extracts the full surplus of the trader by making an offer equal to the trader's value, unless such value is lower than the continuation value R^I. Formally, the price asked by an informed farmer is $p^I(v) = max(v, R^I)$ which is accepted if and only if $v \geq R^I$. It follows that the expected value of being the proposer for an informed farmer is:

$$O^I = E_v p^I(v) = (1 - F(R^I))\frac{v_H + max(R^I, v_L)}{2} + F(R^I)R^I$$

where the expectation is taken over the possible values of v.

Uninformed farmer's strategy

The price offer of the uninformed farmer solves the following maximization:

$$Max_p \quad \int_p^{\infty} pf\left(v\right)dv + R^U \int_{-\infty}^{p} f\left(v\right)dv$$

The first order condition is:

$$\int_{p}^{\infty} f\left(v\right) dv - pf\left(p\right) + R^{U} f\left(p\right) = 0.$$

Under the assumption that v is normally distributed the second order condition is $-2f\left(p\right) < 0$ and the interior solution p^{int} is given by:

$$v_H - p - p + R^U = 0 \quad \Rightarrow \quad p^{int} = \frac{v_H + R^U}{2}.$$

The interior solution is accepted only if $v \geq p^{int}$. It follows that the expected utility of offering price p^{int} is:

$$u^{int} = \left(1 - F(p^{int})\right) p^{int} + F\left(p^{int}\right) R^U$$

Notice that an interior solution exists if $R^U > 2v_L - v_H$, a condition that will turn out useful in the proof of the main proposition. The uninformed farmer can also implement a corner solution, which is offering the price v_L. Such offer is accepted by traders with any value $v \in [v_L, v_H]$ and gives to the farmer a utility equal to v_L. It follows that the optimal strategy of the uninformed farmer is given by:

$$p^U = \begin{cases} p^{int} & \text{if} \quad u^{int} \geq v_L \\ v_L & \text{if} \quad u^{int} < v_L \end{cases}$$

The expected value of being the proposer for the uninformed farmer is:

$$O^U = \begin{cases} (1 - F(p^{int}))p^{int} + F(p^{int})R^U & \text{if} \quad u^{int} \geq v_L \\ v_L & \text{if} \quad u^{int} < v_L \end{cases}$$

where again the expectation is taken over the possible values of v.

Trader's strategy

The trader does not know the farmer's type but only the probability d with which the farmer is informed. Define $\bar{R} = max\{R^I, R^U\}$ and $\underline{R} = min\{R^I, R^U\}$. There are only two possibile optimal strategies: a pooling strategy where the highest continuation value $\bar{R}$ is offered and all types accept, and a separating strategy in which the lowest continuation value is offered and accepted only by one type of farmers. It is easy to see that any other strategy is never optimal: any price $p \in (\underline{R}, \bar{R})$ is only accepted by one type and delivers strictly smaller payoffs than $\underline{R}$. Any price $p > \bar{R}$ is accepted by all farmers and delivers strictly smaller payoffs than offering $\bar{R}$. We can now state the main proposition.

Proposition 1. *There is a threshold probability d^* which is such that (i) for any $d \geq d^*$ the trader's optimal offer is R^I and the offer is accepeted by both types of farmers; and (ii) for any $d < d^*$ the trader's optimal offer is R^U and the offer is only accepted by the uninformed farmer.*

Proof. The proof is divided into 4 lemmas.

Lemma 1. *The continuation value of the informed farmer is at least as high as the continuation value of the uninformed farmer, that is $R^I \geq R^U$.*

Lemma 1 is self-evident. As respondents, informed and uninformed farmers face the same price offer because the trader cannot distinguish between the two types. As proposer, the informed farmer can always mimic the strategy of the uninformed farmer and achieve payoffs that are at least as high.

Lemma 2. *The expected utility of being the proposer for the informed farmer is strictly higher than that of the uninformed farmer, that is $O^I > O^U$.*

We discuss two cases.

Case 1: The uninformed farmer's optimal strategy is the corner solution $p = v_L$.

In this case it is easily verified that

$O^U = v_L < (1 - F(R^I))\frac{v_H + max(R^I, v_L)}{2} + F(R^I)R^I = O^I$

Case 2: The uninformed farmer's optimal strategy is the interior solution $p = \frac{v_H + R^U}{2}$.

We can rewrite O^U as a function of R^U:

$$O^U(R^U) = (1 - F(\frac{v_H + R^U}{2}))\frac{v_H + R^U}{2} + F(\frac{v_H + R^U}{2})R^U = \frac{v_H^2 + R^{U2} + 2R^U v_H - 4R^U v_L}{4(v_H - v_L)}$$

Notice that, for all $R^U > 2v_L - v_H$, O^U is strictly increasing in R^U[24]. It follows that:

$$\begin{aligned} O^I &= (1 - F(R^I))\frac{v_H + R^I}{2} + F(R^I)R^I \\ &> (1 - F(\frac{v_H + R^I}{2}))\frac{v_H + R^I}{2} + F(\frac{v_H + R^I}{2})R^I \\ &= O^U(R^I) \geq O^U(R^U) \end{aligned}$$

Where the first inequality follows from $\frac{v_H + R_I}{2} > R^I$ and the second inequality follows from Lemma 1 and the fact that $O^U(R^U)$ is strictly increasing in R^U.

Lemma 3. *The continuation value of the informed farmer is strictly higher than the continuation value of the uninformed farmer, that is $R^I > R^U$.*

Assume $R^I = R^U = \bar{R}$. Then the optimal strategy for the trader is offering a price equal to $\bar{R}$ which is accepted by both types of farmers.

We can solve for $\bar{R}$ using, for example, the continuation value of the informed farmer:

$$\begin{aligned} \bar{R} &= \beta w O^I + \beta(1 - w)\bar{R} \\ \Rightarrow \bar{R} &= \frac{\beta w O^I}{1 - \beta + \beta w} \end{aligned}$$

It follows that we can write the continuation values as:

$$R^I = \beta w O^I + \beta(1 - w)\bar{R}$$

$$R^U = \beta w O^U + \beta(1 - w)\bar{R}$$

and if we subtract the two equations above from one another we get:

$$R^I - R^U = \beta w(O^I - O^U) > 0$$

[24]Remember that $R^U > 2v_L - v_H$ is a necessary condition for the existence of an interior solution.

Where the inequality follows from Lemma 2, leading to a contradiction.

Lemma 4. *There exists a probability $d^* \in [0,1]$ such that the optimal strategy of the trader is the offer of p^T defined as follows:*

$$p^T = \begin{cases} R^U & \text{if} \quad d \leq d^* \\ R^I & \text{if} \quad d > d^* \end{cases}$$

From Lemma 3 we know $R^I > R^U$. As already shown in the trader's strategy section, the trader chooses between two possible actions. The first is a pooling offer R^I which is accepted by all farmers and gives expected payoffs equal to $v - R^I$. The second is a screening offer R^U which is accepted only by the uninformed farmer and delivers expected payoff $(1-d)(v-R^U)$. The pooling offer dominates the screening offer if and only if:

$$(1-d)(v-R^U) \leq v - R^I$$
$$\Rightarrow d \geq \frac{R^I - R^U}{v - R^U}$$

Hence the threshold strategy defined in Lemma 4 is optimal with $d^* = \frac{R^I - R^U}{v - R^U}$.

□

Bargaining spillovers emerge from traders implementing a pooling equilibrium when they are the proposers. It follows that they only occur when $w \in (0,1)$, that is in situations where both farmers and traders have some bargaining power.[25]

6.2 Further evidence in support of the bargaining model

The model we outlined in the previous section predicts that informed farmers will on average receive higher prices, which is a strong finding in our data. In the model, the adoption of a pooling equilibrium by traders is the mechanism which explains our second empirical finding that, as time goes by, uninformed farmers located in treatment-dense areas also start getting higher prices. These two predictions of the bargaining spillovers model are also consistent with a model of informational spillovers.

[25] If farmers had all the bargaining power, traders are never in the role of proposers. To the contrary, if traders had all the bargaining power, the continuation value of informed and uninformed farmers would be the same and there would not be any difference between pooling and separating equilibria.

However, the model also predicts that informed farmers with $d \leq d^*$ reject the (separating) offers made by traders, while treatment farmers with high values of $d > d^*$ accept the (pooling) offers made by traders[26]. In the data, this could be associated with lower total sales by treatment farmers with low C2T (our proxy for d), and/or with a delay in sales among treatment farmers with low C2T (i.e., a lower fraction of sales made earlier in the agricultural year, due to greater frequency of rejections in earlier months).

This prediction is specific of bargaining spillovers and not consistent with informational spillovers and hence allows us to test our theory. In order to look for evidence of this phenomenon in our data, we compute the cumulative fraction F_{ijt} of yam sold at each point t in the agricultural year for each farmer i[27] and run the following regression:

$$F_{ijt} = \alpha_m + \beta_1 T_j + \beta_2 C2T_j * C_j + \beta_3 C2T_j * T_j + \alpha_s + e_{ijt}, \tag{5}$$

where α_m are monthly fixed effects, and α_s are strata fixed-effects. The model predicts that β_2 is not significantly different from zero and that β_3 is positive. That is, in the control group there should be no relationship between C2T and the cumulative fraction of yam sold at a given point in time. But in the treatment group, at each period in the agricultural year, treatment farmers with higher values of C2T should have sold more of their yam than treatment farmers with lower values of C2T.

The results of this regression for the 2012 and 2013 agricultural years are respectively shown in the first and second columns of Table VIII. In 2012, none of the estimated coefficients are significantly different from zero. However, in 2013, the estimate of β_3 is positive and significant at the 10% level. This supports the model's prediction that treatment farmers with high C2T sell more of their crop early in the season (because of fewer delays in trade).

Figure IV plots the raw means and 90% confidence intervals for the cumulative fraction sold by farmers above and below the sample median of C2T, by month and treatment status. The

[26] Note that, for control farmers, the relationship between d and the probability of a trade occurring is almost always zero. In the few instances where there is a relationship, it goes in the opposite direction for control farmers. The offer made by the uninformed trader is weakly increasing in d (because R^U is weakly increasing in d). There are some cases where for a given v, as d increases, traders start to reject the offer made by the uninformed farmer. So higher values of d reduce the likelihood of trade occurring.

[27] For October 2012 (the third month in the 2012-2013 agricultural season), for example, we calculate the cumulative fraction as the total amount of yam sold from the start of the season through October 2012, divided by the total amount of yam sold over the entire year.

top panel shows results for treatment farmers. In this panel, mean cumulative fraction sold for the above median C2T group is always above the mean for the below median C2T group, and in several months the difference in means is statistically significant. In contrast, the bottom panel shows the results for control farmers. Here, the mean for the high C2T group starts above the mean for the low C2T group, but then changes position several months into the year.

In summary, we derived a simple model of bargaining showing how the presence of network externalities can change the bargaining outcomes for both informed and uninformed farmers. We tested the model's predictions and we found them to be broadly supported by the data. Combining these results with the empirical evidence that fading-out and informational spillovers cannot explain the dynamics of the treatment effect over time, we conclude that the bargaining spillovers are the main driving force. Taking this into account we can now derive de-biased estimates for the treatment effect.

7 Estimating the de-biased treatment effect

Given the spillovers on control group prices, the SUTVA assumption is violated and the estimates of the treatment effect presented in Section 4 are biased. To see this, consider what is being estimated by equation (2), which is repeated here for convenience:

$$p_{ijt} = \sum_{s=0}^{2}\{\lambda_s Y_s + \kappa_s (T_j * Y_s)\} + X'_{ijt}\psi + \alpha_k + \alpha_t + e_{ijt}$$

In this regression, λ_s measures the average control group price in period s, and κ_s measures the difference between average control group price and average treatment group price in period s (controlling for X_{ijt}). Without spillovers, λ_s is an accurate measure of the counterfactual of interest: what treatment group prices would have been absent the intervention. However, when spillovers affect control group outcomes, λ_s no longer represents the counterfactual of interest. Instead, it represents the average control group price inclusive of spillovers. If spillovers cause increases in control group outcomes, then λ_s is biased upward relative to the counterfactual of interest, and κ_s is biased downward relative to the true treatment effect.

In order to de-bias κ_s, we need to find a more accurate measure of the counterfactual of interest.

Ideally, we could use data for a set of "pure control" farmers that we know are completely unaffected by the intervention. However, in our case, we do not have data for a pure control group. Instead, we adapt the techniques developed by Baird, Bohren, McIntosh, and Özler (2014) to generate an estimate of what prices for this pure control group would be.

We make two assumptions to back an estimate of prices for a pure control group: (1) we assume a linear relationship between C2T and prices; and (2) we assume that villages where C2T is equal to zero are unaffected by spillovers. Given these assumptions, we can recover the de-biased treatment effect from our estimates of equation (2) and equation (4), which is reproduced here for convenience:

$$p_{ijt} = \sum_{s=0}^{2} \{\delta_s Y_s + \alpha_s (T_j * Y_s) + \beta_s (C_j * C2T_j * Y_s) + \gamma_s (T_j * C2T_j * Y_s)\} + \alpha_s + \alpha_t + X'_{ijt}\psi + e_{ijt}$$

Equation (4) is similar to (2), but it also includes $C2T_j$ interacted with treatment status. In this equation, δ_s is equivalent to $E[p_{ijt}|T_j = 0, C2T_j = 0]$ in period s. In other words, it is a measure of mean prices for a pure control group unaffected by spillovers (given our two assumptions). The average spillover effect for the control group is equal to the difference between the observed average price in the control group (λ_s) in equation (2) and the estimated pure control average (δ_s) in equation (4). To de-bias κ_s, we need to net out the average spillover effect on the control group. So, the unbiased treatment effect is equal to $\kappa_s + (\lambda_s - \delta_s)$, i.e. the biased treatment effect adjusted for the impact of spillovers on the control group.[28]

Table IX shows the estimates of (2) and (4) using the monthly sales data.[29] Table X presents estimates of the de-biased treatment effect for each time period, as well as estimates of the average spillover effect on control group farmers. The de-biased treatment effect is estimated to be 16.50 GHS in Year 1, and 17.62 GHS in Year 2, significant at the 10% level. These estimates are substantially higher than the biased treatment effect estimates presented in Section 4, due to large positive spillovers on control group farmers. Note that these figures are on-par with the estimated treatment effect in the first few months of the intervention, as shown in Figure I.

Compared to many other interventions in the agricultural marketing space, our estimated treat-

[28] We can also use this approach to estimate the average spillover effect on the treatment group. This is equal to the de-biased treatment effect less α_s, which is the treatment effect for a farmer with a C2T score of zero (which, by our assumption, implies that the farmer is completely unaffected by spillover effects associated with the intervention).

[29] We estimate the two equations using two-step GMM so that we can conduct significance testing on our estimates of average spillover effects and the de-biased treatment effect.

ment effects are fairly substantial. The estimates of the de-biased treatment effects represent 13.1% of the estimated mean price for the pure control group in Year 1, and 11.4% of average pure control group price in Year 2.[30] For the average yam farmer selling 1,200 tubers in a year, the treatment effect translates into an additional 200 GHS (US$135) in annual revenue.[31] Considering that profit margins for farmers are believed to be low, the impact could be considerably larger in terms of an increase in farm profits. We did not collect information on farmers' costs, so we are unable to provide an exact figure for the impact in terms of profits. However, assuming a profit margin of 50% and no change in costs an 11%-13% increase in prices would translate into a 22%-26% increase in profits.

Another way to consider the magnitude of the intervention is to compare the cost of the service with the estimated benefits to farmers. Esoko recently started offering an annual subscription to farmers for 24 GHS (18 GHS in real August 2011 cedis), which is comparable to the per-farmer cost that we paid Esoko in our study. Ignoring the costs associated with training, the de-biased treatment effect for Year 1 implies a return on investment that is about ten times larger than the cost of the service for the average yam farmer. Adding in the costs of training, the estimated ROI is over 200%.[32]

8 Conclusion

In this paper, we study the impact of providing better price information to rural farmers in Ghana. We show that the alerts had a large and statistically significant impact on producer prices for yam in the first year of the study, but that this effect disappeared in the longer term. We show that this is neither due to a fade way of salience nor to information sharing and we propose a "bargaining spillovers" explanation that is largely consistent with our data.

Our exploration of the reasons behind the disappearance of the treatment effect is critical to interpreting our long-term results. A spillover story actually implies that uninformed farmers in-

[30] As a point of comparison, Goyal (2010) finds that information kiosks introduced by ITC (a major processor) in Madhya Pradesh, India, increased market prices for soybeans by 1-3%. Many other MIS evaluations have failed to find any effect on prices, while a few (Nakasone, 2013; Svennson and Yanagizawa, 2009) have estimated effects that are closer to 15%.

[31] GHS figure in real Ghana cedis as of August 2011. Exchange rate calculated using historical data from Oanda.com for August 2011.

[32] Our costs of training were about US$37 per farmer for our study, or about GHS 60 in real August 2011 cedis.

directly benefit from the MIS, which means that our estimate of the long-run average treatment effect underestimates (potentially severely) the impact of the service on farmers. More generally, these results show that indirect impacts of interventions in rural agricultural markets can be substantial, can have important effects on individual-level treatment effect estimates, and are therefore extremely important to take into consideration.

A final point worth discussing is the impact that the price alerts had on farmers' marketing behaviors. In the Appendix, we show that the alerts caused some farmers to start selling in urban markets, and caused some farmers to delay sales of their yams until later in the marketing season, a finding which is consistent with our model of bargaining spillovers. While changes along these dimensions of marketing behavior were quite small in aggregate during the first year of the study, we believe that, with even more time, farmers may actually start to change where they sell in response to market information. They may even change their production decisions. We leave the exploration of this subject to future research.

References

[1] Aker, J. (2008), “Does Digitial Divide or Provide? The Impact of Cell Phones on Grain Markets in Niger,” mimeo, UC Berkeley.

[2] Aker, J. and M. Fafchamps (2010), “How Does Mobile Phone Coverage Affect Farm-Gate Prices? Evidence from West Africa,” mimeo, December.

[3] Asunka, J., S. Brierley, M. Golden, E. Kramon, and G. Ofosu (2014), “Protecting the Polls: The Effect of Observers on Election Fraud,” mimeo (February).

[4] Baird, S., J.A. Bohren, C. McIntosh, and B. Özler (2014), “Designing Experiments to Measure Spillover Effects,” mimeo (February 7).

[5] Bruhn, M. and D. McKenzie (2009), “In Pursuit of Balance: Randomization in Practice in Development Field Experiments,” *American Economic Journal: Applied Economics* 1(4): 200-232.

[6] Camacho, A. and E. Conover (2011), “The Impact of Receiving Price and Climate Information in the Agricultural Sector,” mimeo, November.

[7] Crépon, B., E. Duflo, M. Gurgand, R. Rathelot, and P. Zamora (2013), ”Do Labor Market Policies Have Displacement Effects? Evidence from A Clustered Randomized Experiment,” *Quarterly Journal of Economics* 128 (2): 531-580.

[8] Cudjoe, G., C. Breisinger, and X. Diao (2010), “Local Impacts of a Global Crisis: Food Price Transmission, Consumer Welfare, and Poverty in Ghana,” *Food Policy* 35: 294-302.

[9] Duflo, E. R. Glennerster, and M. Kremer (2007), “Using Randomization in Development Economics Research: A Toolkit,” in *Handbook of Development Economics*, T.P. Schultz and J.A. Strauss, eds, Vol. 4, Chapter 61, Elsevier: 3895-3962.

[10] Fafchamps, M. and B. Minten (2012), “Impact of SMS-Based Agricultural Information on Indian Farmers,” *World Bank Economic Review.*

[11] Giné, X. and G. Mansuri (2012), “Together We Will: Experimental Evidence on Female Voting Behavior in Pakistan,” mimeo, August.

[12] Goyal, A. (2010), "Information, Direct Access to Farmers, and Rural Market Performance in Central India," *American Economic Journal: Applied Economics* 2 (July): 22-45.

[13] Jensen, R. (2007), "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector," *Quarterly Journal of Economics* 127 (3): 879-924.

[14] Fafchamps, M. and B. Minten (2012), "Impact of SMS-Based Agricultural Information on Indian Farmers," *World Bank Economic Review.*

[15] Ministry of Food and Agriculture (2010), "Agriculture in Ghana: Facts and Figures (2009)," December.

[16] Miguel, E. and M. Kremer, "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica* 72 (1): 159-217.

[17] Mitra, S., D. Mukherjee, M. Torero, and S. Visaria (2013) "Asymmetric Information and Middlemen Margins: An Experiment with West Bengal Potato Farmers," mimeo, August.

[18] Muto, M. and T. Yamano (2009) "The Impact of Mobile Phone Coverage Expansion on Market Participation: Panel Data Evidence from Uganda," *World Development* 37 (2): 1887-1896.

[19] Myerson, R. (1984), "Two-Person Bargaining Problems with Incomplete Information," *Econometrica* 52: 461-487.

[20] Nakasone, Eduardo (2013) "The Role of Price Information in Agricultural Markets: Experimental Evidence from Rural Peru," Job Market Paper, November.

[21] Svensson, J. and D. Yanagizawa (2009), "Getting Prices Right: The Impact of the Market Information Service in Uganda," *Journal of the European Economic Association* 7 (2-3): 435-445.

[22] Svensson, J. and D. Yanagizawa-Drott (2012), "Estimating Impact in Partial vs General Equilibrium: A Cautionary Tale from a Natural Experiment in Uganda," mimeo (August).

Table I: Background on agricultural marketing, by crop

	Yam	Maize	Cassava	Groundnut
Where crop is sold				
Percent sell at farm-gate or home	55.6%	71.1%	77.6%	54.2%
Percent sell at local markets	61.5%	49.9%	64.8%	60.0%
Percent sell at urban markets	24.0%	1.5%	0.5%	1.3 %
Main place of sale				
Farm-gate or home	39.2%	58.0%	54.4%	44.4%
Local markets	46.5%	41.5%	45.4%	55.1%
Urban markets	15.3%	0.2%	0.2%	0.9 %
Bargaining				
Percent that bargain with traders	99.6%	51.7%	40.5%	26.1%
Long-term relationships with buyers				
Percent with any long-term buyers	69.9%	70.6%	87.3%	64.9%
Mean number of long-term buyers	3.78	3.62	5.05	3.71
Median number of long-term buyers	2	2	3	2

Data come from the baseline survey, except for information on bargaining which comes from the midline survey. Mean and median number of long term buyers are among those farmers with at least one long-term buyer.

Table II: Descriptive statistics and balance, baseline data

	Control	Treatment	T - C
Farmer characteristics			
Age	41.00	40.60	-0.40
Schooling - JHS or higher	44.7%	38.8%	-5.9%
Male	78.7%	81.7%	2.9%
Farming is main source of income	76.8%	79.7%	2.8%
Land cultivated last season (acres)	6.72	7.21	0.50
Median income from two main crops (GHS)	1,400	1,400	0
Mean income from two main crops (GHS)	2,064	2,320	256
Mean of asset index	0.081	-0.077	-0.158
Owns a bicycle	83.2%	82.2%	-0.9%
Owns a motorbike	27.7%	29.8%	2.0%
Owns a radio	73.1%	71.2%	-1.9%
Owns a TV	36.4%	30.4%	-6.1%
Phone ownership and usage			
Owns a mobile phone	72.3%	69.8%	-2.5%
Sends SMS messages	22.6%	14.7%	-7.9%*
Receives SMS messages	32.0%	22.9%	-9.1%
Crops grown			
Yam	60.7%	65.9%	5.1%
Cassava	37.0%	43.8%	6.8%
Maize	46.1%	35.7%	-10.4%*
Groundnut	19.2%	26.0%	6.8%
Where crops are sold			
Percent sell at farm/home	73.6%	75.1%	1.6%
Percent sell at local markets	67.6%	65.7%	-1.9%
Percent sell at urban markets	15.5%	18.7%	3.2%
Mean distance to nearest district market (km)	10.71	11.07	0.36
Knowledge of market prices			
Percent well informed about urban prices	33.3%	26.2%	-7.2%
Percent well informed about local prices	84.6%	75.1%	-9.5%
Number of communities	49	51	
Number of clusters	45	45	

Standard errors of the difference are clustered at the community level.
"Sends SMS" and "Receives SMS" figures include mobile phone owners only.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

Table III

Impact of price alerts on yam prices, assuming no spillovers

	Year 1		Year 2		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment, Pre-T					-0.099	1.125
					(6.945)	(6.938)
Treatment, Year 1	6.256*	7.707**			7.332*	8.693**
	(3.727)	(3.454)			(3.838)	(3.759)
Treatment, Year 2			0.474	1.185	-0.202	0.433
			(4.497)	(4.581)	(4.507)	(4.549)
N	1519	1519	2654	2654	5020	5020
R^2	0.294	0.300	0.221	0.224	0.311	0.314
Control group mean	134.15	134.15	172.07	172.07	150.83	150.83
Other covariates		✓		✓		✓

Notes: Prices are per 100 tubers, denominated in real, August 2011 Ghana Cedis (GHS). Year 1 = Nov 2011-Jun 2012. Year 2 = Jul 2012-Jun 2013. Pre-T = Aug 2011-Oct 2011 (before the start of the intervention). All regressions include strata fixed effects, period fixed effects, and controls for yam type. Other covariates include farmer's gender and asset index level, and the community's distance to the closest district market. Standard errors clustered at the community cluster level are shown in parentheses. Analysis relies on monthly data; results using annual data are comparable and are available upon request.
** Significant at 5% level. * Significant at 10% level.

Figure I

Impact of price alerts on yam prices, assuming no spillovers

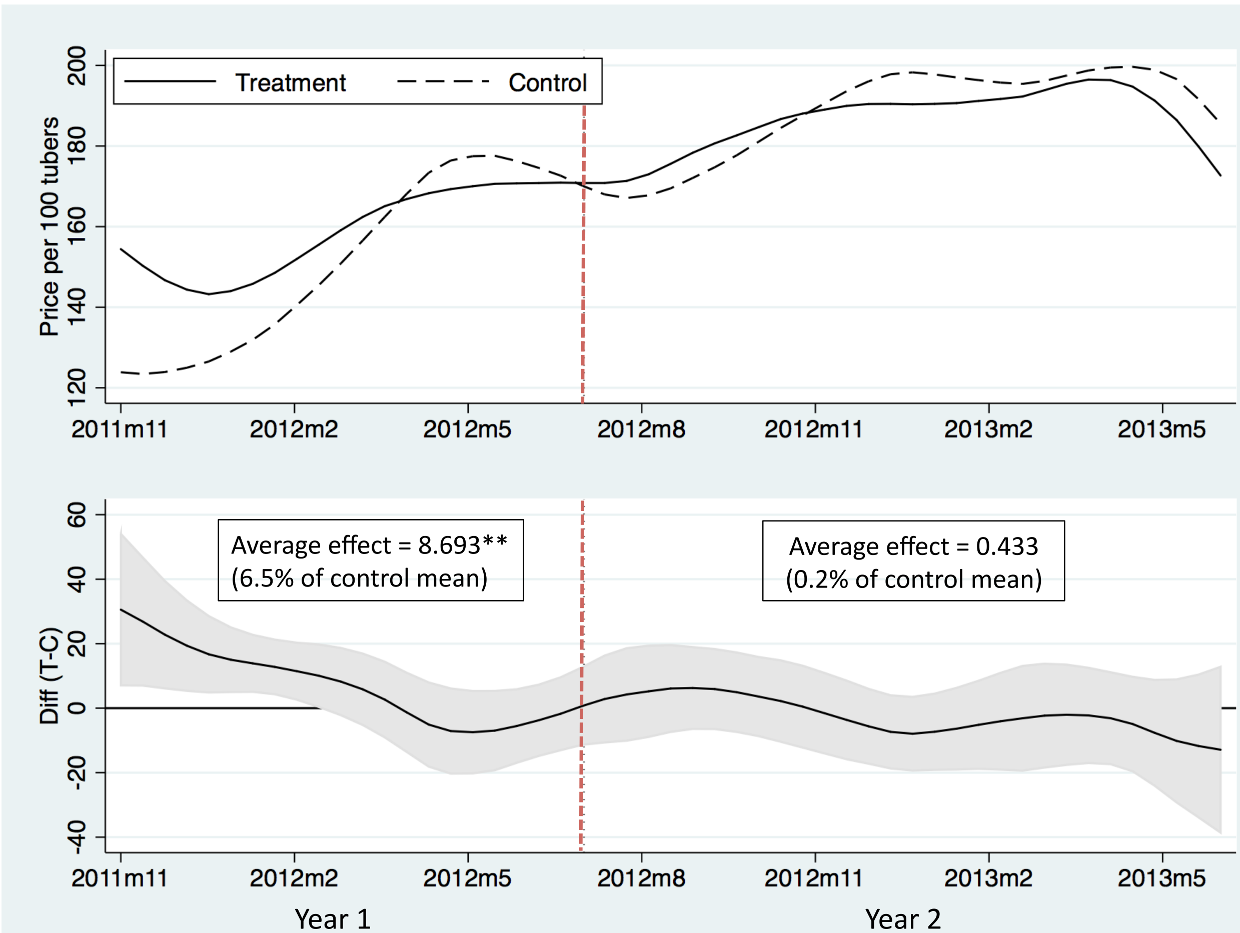

Notes: The top figure plots yam prices for treatment and control groups, estimated using fan regressions (controlling for strata fixed effects, yam type, gender, asset index, and the community's distance to the nearest local market). The bottom figure plots the difference between treatment and control group prices, with the bootstrapped 95% confidence interval shown in grey (block bootstrap at the community cluster level, 1000 replications with replacement). The bottom figure also displays the average estimated treatment effect for each agricultural year, using results from the pooled regression with additional covariates (column (6) of Table III).

Table IV: Farmers self-reported views on Esoko, Annual surveys

	Midline	Endline
Farmers receiving Esoko		
Treatment	89.1%	90.9%
Control	1.7%	0.2%
If receive: are the alerts useful?		
Yes	78.9%	92.5%
No	1.6%	0.7%
Need more time to decide	19.5%	6.8%
If receive: have you used the alerts to...		
Bargain with buyers?	67.5%	83.5%
Decide where to sell?	37.9%	68.3%
Decide when to sell?	22.0%	52.3%
Make production decisions?	11.2%	19.0%
If receive: how have the alerts affected your prices?		
Higher	53.0%	79.0%
No change	15.4%	2.7%
Lower	0.0%	0.7%
Don't know	10.5%	1.4%
Have not used the alerts	20.8%	16.3%
If higher prices: what is main reason?		
I could select the best time to sell	18.5%	19.5%
I could select the best market to sell in	18.5%	8.9%
I have better bargaining power now	63.0%	71.4%

Figure II: Impact of C2T on yam prices over time

(a) **Estimated coefficients (β_s and γ_s)**

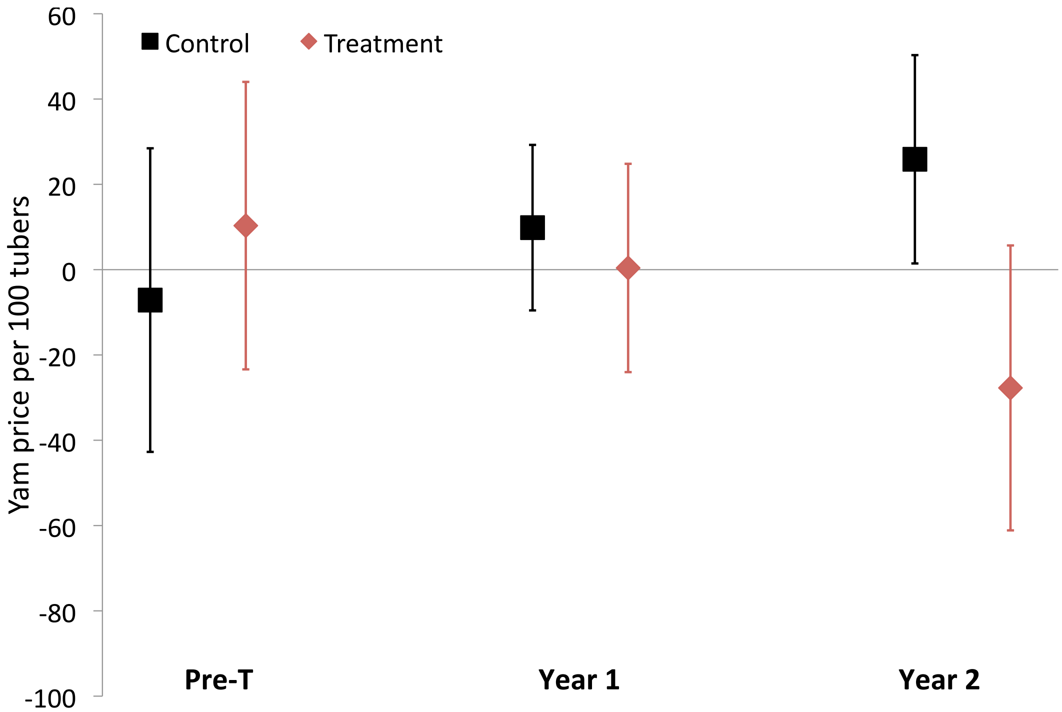

(b) **Difference (β_s - γ_s)**

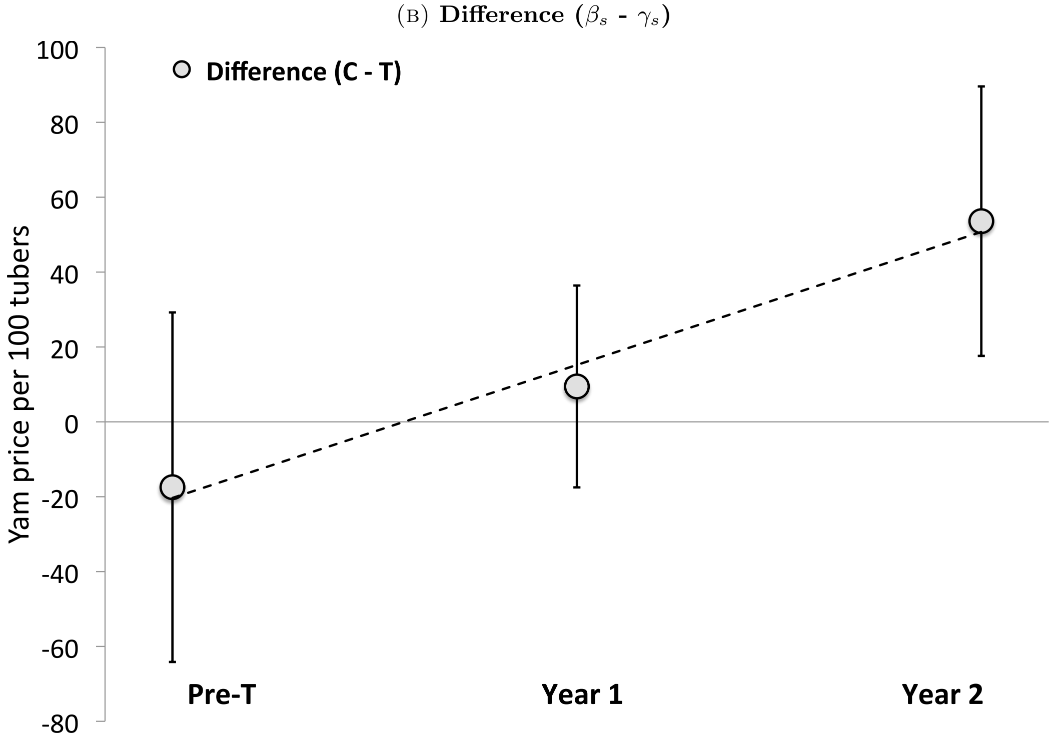

Notes: The top panel plots β_s (impact of C2T on prices for the control group) and γ_s (impact of C2T on prices for the treatment group). The bottom panel shows the difference in the impact of C2T on prices ($\beta_s - \gamma_s$). The error bars show 90% confidence intervals. The dashed line in the bottom panel is a linear trend line.

Figure III: Impact of C2C on yam prices over time (placebo test)

(A) **Estimated coefficients (β_s and γ_s)**

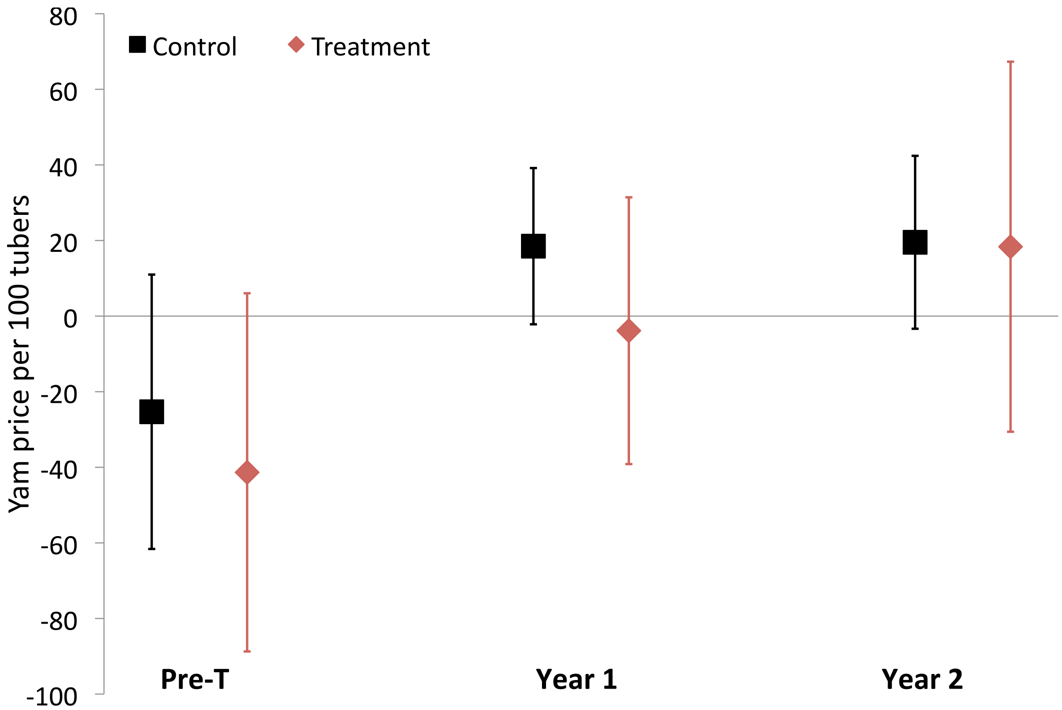

(B) **Difference (β_s - γ_s)**

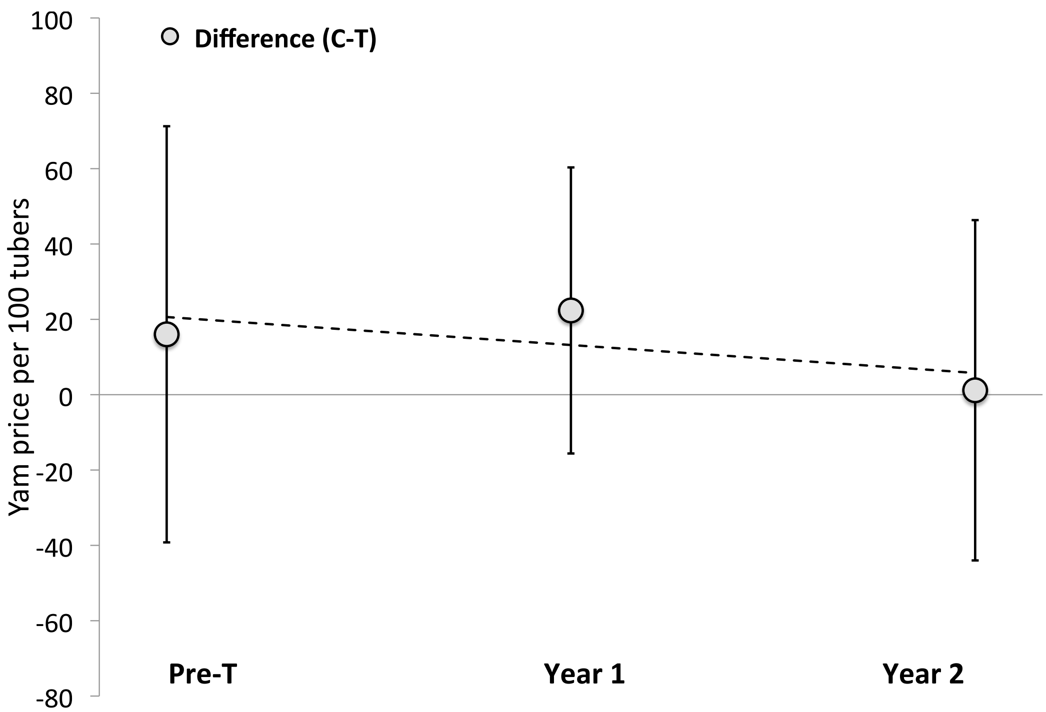

Notes: The top panel plots β_s (impact of C2C on prices for the control group) and γ_s (impact of C2C on prices for the treatment group). The bottom panel shows the difference in the impact of C2C on prices ($\beta_s - \gamma_s$). The error bars show 90% confidence intervals. The dashed line in the bottom panel is a linear trend line.

Table V: Farmers' feelings about being "well informed" about market prices

	Baseline		Endline	
	(1)	(2)	(1)	(2)
Panel A: Urban markets				
Treatment	-0.458**	-0.592	0.822***	0.778*
	(0.189)	(0.459)	(0.178)	(0.455)
C2T * Control		-0.295		-0.715
		(0.667)		(0.717)
C2T * Treatment		-0.061		-0.687
		(0.723)		(0.669)
N	609	609	622	622
Panel B: Local markets				
Treatment	-0.574**	0.555	0.487***	1.552***
	(0.287)	(0.690)	(0.174)	(0.411)
C2T * Control		2.206**		0.423
		(0.910)		(0.741)
C2T * Treatment		-0.183		-1.686***
		(1.272)		(0.577)
N	616	616	622	622

Farmers were asked to respond to the following questions: "Do you feel that you are well informed about URBAN [LOCAL] market prices?" We present results from ordered probit regressions, where answers are coded as: 1 = "no, not at all"; 2 = "no, not very well"; 3 = "yes, fairly well"; 4 = "yes, very much". All regressions include strata fixed-effects. In the table above, we only include farmers that sell yam. Huber-White robust standard errors clustered by community cluster are in parentheses.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

Table VI: Number of markets for which farmers regularly know market prices

	Baseline		Midline	
	(1)	(2)	(1)	(2)
Treatment	-0.117	0.309	0.514***	0.751***
	(0.120)	(0.398)	(0.101)	(0.289)
C2T * Control		0.853*		0.496
		(0.501)		(0.497)
C2T * Treatment		0.146		0.079
		(0.552)		(0.287)
Control group mean	1.70	1.70	1.33	1.33
Treatment group mean	1.60	1.60	2.17	2.17
N	628	628	634	634

We asked farmers to list the number of markets (local and urban) for which they regularly knew market prices. We took these lists and generated counts of the number of markets listed by each farmer (ranging from 0 to a maximum of 4). We present results from a Poisson (count data) regression. All regressions include strata fixed-effects. In the table above, we only include farmers that sell yam. Huber-White robust standard errors clustered by community cluster are in parentheses.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

Table VII: Estimation errors of yam prices in Accra, endline survey

	Log of absolute error		Log of absolute % error	
	(1)	(2)	(1)	(2)
Treatment	-0.244*	-0.554	-0.282*	-0.580
	(0.141)	(0.349)	(0.159)	(0.381)
C2T * Control		-0.298		-0.291
		(0.474)		(0.529)
C2T * Treatment		0.309		0.295
		(0.623)		(0.698)
N	541	541	541	541
R^2	0.103	0.104	0.095	0.096

The endline survey asked farmers to provide an estimate of contemporaneous prices for yam in Accra. We calculated "errors" by taking the difference between the price provided in the Esoko alerts and the farmer's estimate. All regressions include strata fixed effects, interview week fixed effects, and yam type fixed effects. Huber-White robust standard errors clustered by community cluster are in parentheses.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

Table VIII

Impact of C2T on timing of yam sales

	Year 1	Year 2
Treatment	0.031	-0.106*
	(0.066)	(0.063)
C2T * Control	0.100	0.031
	(0.082)	(0.042)
C2T * Treatment	0.092	0.208*
	(0.121)	(0.105)
N	4490	6831
R^2	0.571	0.580

Notes: The dependent variable is the cumulative fraction of yam that each farmer has sold by a given period in the agricultural year. The regression includes monthly fixed effects and strata fixed effects. Standard errors are adjusted for clustering at the community cluster level. * Significant at 10% level

Figure IV: Cumulative fraction of yam sold in AY2013, by month and C2T level

(a) Treatment

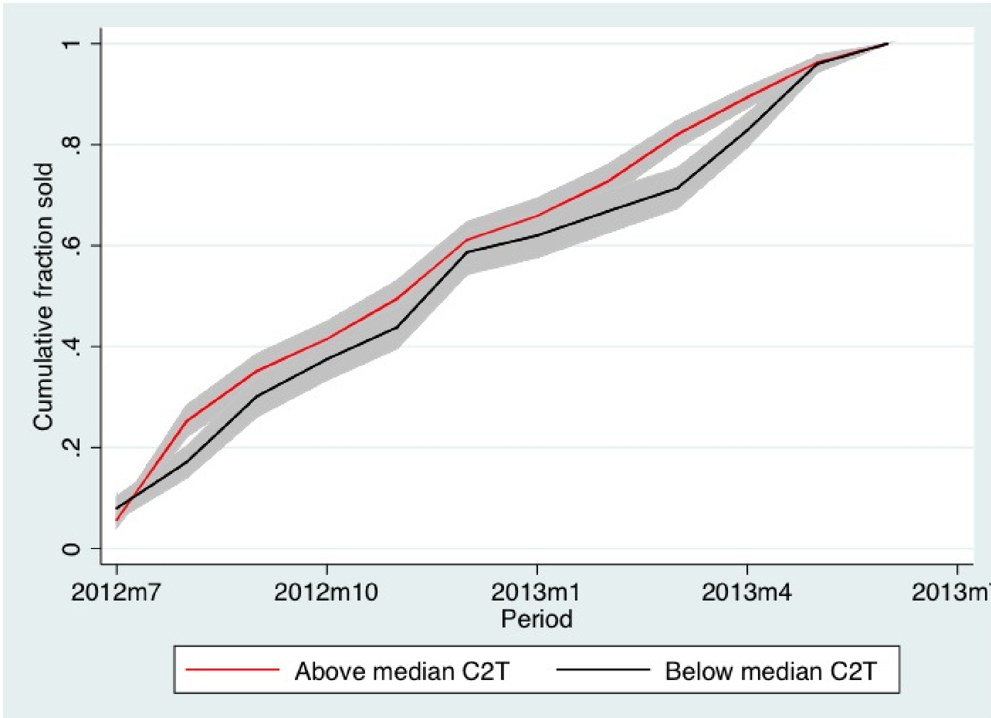

(b) Control

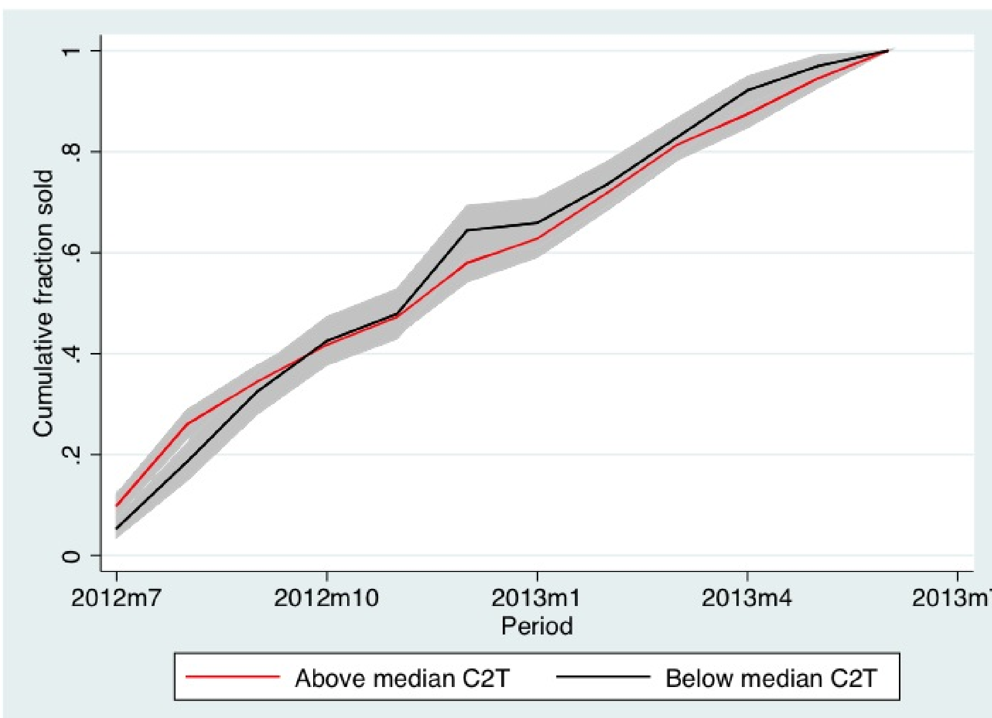

Notes: The figures plot the mean cumulative fraction sold by each month, for farmers with above median C2T and for farmers with below median C2T. 90% confidence intervals are shaded in grey.

Table IX

Estimating the de-biased treatment effect

	Equation (2)	Equation (4)
Treatment, pre-T	1.125	-9.026
	(6.868)	(16.003)
Treatment, Year 1	8.693**	13.854
	(3.721)	(9.730)
Treatment, Year 2	0.433	29.839***
	(4.503)	(11.443)
C2T * Control, Pre-T		-7.139
		(21.163)
C2T * Control, Year 1		9.847
		(11.536)
C2T * Control, Year 2		25.875*
		(14.522)
C2T * Treatment, Pre-T		10.319
		(20.038)
C2T * Treatment, Year 1		0.395
		(14.519)
C2T * Treatment, Year 2		-27.730
		(19.862)
Pre-T	98.322***	101.034***
	(9.615)	(16.262)
Year 1	158.123***	150.321***
	(12.350)	(15.441)
Year 2	163.114***	145.931***
	(11.335)	(16.116)

Notes: Prices are per 100 tubers, denominated in real, August 2011 Ghana Cedis (GHS). Year 1 = Nov 2011-Jun 2012. Year 2 = Jul 2012-Jun 2013. Pre-T = Aug 2011-Oct 2011 (before the start of the intervention). Regressions include strata fixed effects, period fixed effects, controls for yam type, and other covariates (farmer's gender and asset index level, and the community's distance to the closest district market). Standard errors clustered at the community cluster level are shown in parentheses. The equations are estimating using two-step system GMM. *** Significant at 1% level ** Significant at 5% level * Significant at 10% level

Table X

Estimate of de-biased ITT

	(1) Pre-T	(2) Year 1	(3) Year 2
Panel A: Price per 100 tubers			
Biased treatment effect (κ_s)	1.125	8.963**	0.433
	[-10.17, 12.42]	[2.84, 15.08]	[-6.97, 7.84]
Average spillovers for Control ($\lambda_s - \delta_s$)	-2.712	7.802	17.183*
	[-24.73, 19.30]	[-5.25, 20.86]	[0.77, 33.59]
De-biased treatment effect [$\kappa_s - (\lambda_s - \delta_s)$]	-1.587	16.495*	17.616*
	[-24.38, 21.21]	[1.97, 31.02]	[0.87, 34.36]
Panel B: As % of pure control price			
Biased treatment effect (κ_s)	0.9%	7.1%**	0.3%
	[-8.4%, 10.3%]	[2.2%, 11.9%]	[-4.5%, 5.1%]
Average spillovers for Control ($\lambda_s - \delta_s$)	-2.3%	6.2%	11.1%*
	[-20.5%, 16.0%]	[-4.2%, 16.5%]	[0.5%, 21.7%]
De-biased treatment effect [$\kappa_s - (\lambda_s - \delta_s)$]	-1.3%	13.1%*	11.4%*
	[-20.2%, 17.6%]	[1.6%, 24.6%]	[0.6%, 22.2%]

Notes: Figures in square brackets denote 90% confidence intervals.
*** Significant at 1% level ** Significant at 5% level * Significant at 10% level

A Formation of community clusters

In this Appendix we provide a detailed description of how we formed community clusters for the randomization. All information used for community cluster formation comes from the baseline survey. We asked farmers for information about where they sold their crops, and who they spoke to about their marketing. We also gathered GPS coordinates for each village. From this information, we constructed three indices of "connectedness" for each village-pair combination.

A.1 Creation of indices

A.1.1 Market overlap index

The market overlap index measures the extent to which farmers in communities i and j overlap in their marketing activities. We asked each farmer to list up to three markets where they had sold their production in the previous agricultural season. We then used this information to identify the number of farmers in a given community that sell in each market. Let n_{im} represent the number of farmers in community i that report selling in market m, and n_{jm} represent the number of farmers in community j that report selling in market m. To come up with a measure of market overlap for communities i and j, we multiply n_{im} and n_{jm} together for each market m, and sum over all the possible markets:

$$mo_{ij} = \sum_{m=1}^{M} n_{im} n_{jm} \tag{6}$$

In this calculation, we ignore overlapping sales in Accra. Given the size of the Accra market, we don't believe it is likely that farmers in our sample would actually encounter one another in the Accra market, or would otherwise be affected by the presence of farmers from other study communities.

A.1.2 Marketing communications index

In the baseline survey, we asked people to list up to two communities that they communicate with about their marketing. Farmers were also asked to provide details on:

- Frequency of communication: daily (which we code =1), weekly (=2), or occasionally (=3).

Note that lower values reflect more frequent communication.

- Number of contacts in the community. The options were: many (=1), few (=2), or one (=3). Lower values reflect more contacts.

Let f_{nij} represent the frequency with which farmer n in community i communicates with people in community j, and c_{nij} represent the number of contacts that farmer n in community i has with people in community j. We take this information an construct a single measure of communication intensity $s_{nij} = 7 - f_{nij} - c_{nij}$, which can range from 1 (lowest intensity) to 5 (highest intensity). We set s_{nij} equal to zero for all communities that are not mentioned by a farmer.

To construct measure of marketing communications between communities i and j, we add together the sum of the s_{nij} for farmers in community i and the sum of the s_{nji} for farmers in community j:

$$mc_{ij} = \sum_{n=1}^{N_i} s_{nij} + \sum_{n=1}^{N_j} s_{nji} \tag{7}$$

A.1.3 Geographic proximity index

Finally, we use GPS coordinates for each community to identify the distance (as-the-crow-flies) between each community pair i and j. In our geographic proximity index, gp_{ij}, we multiply distances (reported in km) by negative 1 so that a larger number represents closer geographic proximity.

A.2 Cluster formation

Once we calculated the three indices described above, we needed to find a way to combine them into a single measure of connectedness, c_{ij}, that we could use for cluster formation. We started by standardizing all indices to have a mean of 0 and standard deviation of 1, i.e. we created z-scores. This was due to the fact that each index is on its own scale, and there is no clear way to meaningfully compare values across the different scales. Next, we ran principal components analysis on the three standardized indices. Principal components is a statistical technique that helps to reduce the dimensionality of data while maintaining as much information as possible. We used the first principal component (which in our case, explains about 53% of the total variance in the data) to calculate a weighted average of our three indices. The weights we used in this calculation were provided by the the first principal component, and were equal to 0.6381 for distance, 0.4565 for marketing communication, and 0.6201 for market overlap:

$$c_{ij} = 0.6381(gp_{ij}) + 0.4565(mc_{ij}) + 0.6201(mo_{ij}) \tag{8}$$

The next step was choosing a cut-off value for c_{ij}, above which communities i and j would be considered connected enough to warrant assignment to the same community cluster, and below which they would be kept in separate clusters. We faced a tradeoff between setting a low cut-off value, which would take a very conservative approach to minimizing spillovers but would reduce the number of units of randomization in our sample, and a high cut-off value, which would allow us to maintain a larger number of randomization units but would potentially leave us more exposed to spillovers. We combined our results for the c_{ij} and the anecdotal information we gathered during our field work to settle on a cut-off value of 6. This value ensured that communities we knew to be highly connected were grouped into the same cluster, but also kept the total number of community clusters large (90 in total).

B Mechanisms behind the short-term average treatment effect

The price alerts had a positive and economically significant effect on yam prices in the first year of the study. In this Appendix, we investigate the channels through which this effect may have occurred, including: (a) improvements in farmers' bargaining position with traders; (b) changes in place of sale; (c) changes in timing of sales; and (d) changes in production decisions. Our data overwhelmingly point to the bargaining channel as being the mechanism through which the alerts improved prices in the first year of the study. This is consistent with our failure to detect a price effect for other types of crops, since bargaining is an essential component of yam marketing, but is far less important or widespread for the other crops that we study.

B.1 Bargaining

One way that the price alerts could have increased producer prices is by resolving informational asymmetries between farmers and traders, giving farmers a better position in negotiations. To determine whether and how the alerts affected farmers' bargaining with traders, we rely on data from the follow-up survey. In that survey, we asked farmers to recall an "important" sales transaction for their main crop from the previous agricultural season.[33] We collected basic information about the sale (where it took place, time of sale, quantity sold, type of buyer) as well as more detailed information about whether bargaining took place, who made the first price offer, and the farmer's price expectations on the day prior to the sale.

In our data, virtually all yam farmers reported (1) bargaining with traders in their yam sales, and (2) making the initial price offer in these negotiations. Table A1 provides a detailed analysis of these negotiations. The table reports OLS regressions of (a) farmer's initial offer price, (b) final price received, and (c) price expectation in the day prior to the sale. We only include sales where bargaining took place and the farmer made the first price offer. Controlling for quantity sold, place of sale, strata, and month of sale, we find that treatment group farmers had a higher expectation of the price they could get for their yam, by almost 50 cedis per 100 tubers (about 20% of the final price received for control group farmers). They made an initial price offer that was about 44 cedis higher than the price offer made by the control group, and ended up getting

[33] The goal was to have the farmer choose a sale that was significant to his income for that year, and that he also remembered well enough to provide answers to our questions.

a price that was about 24 cedis higher than the control group (about 10% of the final price for control group farmers). Overall, the table provides strong evidence that the price alerts changed price expectations, resulting in treatment farmers making higher initial offers in their negotiations with traders, and getting higher final prices as a result.

B.2 Place of sale

Most of the literature on mobile phones and agricultural MIS focuses on the impact of better access to information on farmers' decisions about *where* to sell. The idea is that better market information reduces search costs, enabling farmers to identify and exploit spatial arbitrage opportunities for private gain. To analyze whether the alerts affected farmers' decisions about where to sell, we use information from the annual surveys to create indicator variables for whether a farmer made any sales at the farm gate/home, in local markets, or in urban markets in the previous agricultural season. We then run regressions of each of these indicator variables on the treatment dummy, strata fixed effects, and (in some specifications) additional individual and community controls. To improve the precision of the estimates in the follow-up and endline regressions, we include the pre-treatment outcome as an additional control.

Table A2 shows the results of these regressions for yam. In the first year of the intervention, the price alerts led more treatment group farmers to sell in urban markets (relative to the control group). This result is statistically significant and robust across the two specifications. In the endline results, the treatment effect in Panel A is still positive, but it is lower in magnitude and not significantly different from zero. There is some weaker evidence that the alerts also: (a) caused fewer farmers to sell in local markets in Year 1, and (b) induced more farmers to sell at the farm gate in Year 2. However, both of these latter results lose significance once individual controls are added to the regression.

We have shown that the price alerts increased the proportion of farmers selling yams in urban markets, at least in Year 1. Since Year 1 is also when we find a significant impact on producer prices, one possibility is that the switch to urban markets is the driving force behind our positive treatment effect. This is particularly true because, in the regressions presented so far, we have not included a control for place of sale. For the annual data, it is not possible to do this, since farmers make sales at multiple locations over the year, and we only have aggregate (annual) price information. However, using the monthly transactional data, we can add controls for where crops are sold and see how this affects the estimated treatment effect. Note that place of sale is an endogenous variable, so interpreting these regressions is not straightforward.

Table A3 shows the estimated treatment effect in the first year of the study (through June 2012) using the monthly data, with and without the inclusion of dummy variables for place of sale (local market, urban market, farm gate, or home). The results in columns (1) and (3) omit the indicator variables for place of sale; these are the results that are most comparable to the treatment effect results reported in Table **??** using the annual data. Comparing Table A3 and Table **??**, we see that the estimated ATE for Year 1 is slightly smaller when the monthly data are used, but the estimates are very close to one another.

The estimated treatment effects in columns (2) and (4)—which control for place of sale—are very similar to the estimated treatment effects without controls for place of sale. If movement to urban markets was the principal mechanism through which yam farmers achieved higher prices, then the inclusion of these indicators variables should result in a coefficient on the treatment dummy that is not significantly different from zero. From this, we conclude that changes in place of sale are *not* the main mechanism behind the positive treatment effect in the first year.

B.3 Timing of sales

In addition to affecting decisions about where to sell, the price alerts could have impacted decisions made about the timing of sales over the course of the agricultural season. Providing weekly price information across different markets may help farmers to study price trends over the season, which could result in more profitable decisions about when to sell. This type of effect might be particularly true for crops that exhibit lots of seasonality in prices, such as yam.

We rely on the monthly data to study changes in farmers' selling decisions over time. Arguably the most important dimension of timing that the price alerts could affect is decisions about whether to (a) sell early in the agricultural season, around harvest time, when supply is higher and prices are often lower; or (b) wait to sell later in the agricultural season, in the "lean" season of March to May, when supply is lower and prices tend to be higher. In order to study the impact of the price alerts on this decision, we identify the *last* month in the agricultural season that the farmer made any sales of yam. A farmer that sells everything at harvest time will have made his last sale much earlier in the agricultural season than a farmer that waits to sell some production during the lean season. We plot the empirical cumulative distribution functions (CDFs) of last month of sales separately for treatment and control groups. These CDFs are depicted in Figure A1; the top figure presents results for the 2012 agricultural season (July 2011-June 2012), and the bottom figure presents results for the 2013 agricultural season (July 2012-June 2013). The CDF for the 2012 agricultural season does not appear to differ much between the treatment and control groups. This is confirmed by Kolmogorov-Smirnov (KS) tests for distributional equality: the p-value is very far away from conventional levels of significance. In contrast, the 2013 results indicate statistically significant differences between treatment and control. It appears that treatment group farmers tended to have later final months of sale than the control group; e.g. by March 2013, about 45 percent of control group farmers had sold all of their production for the year, versus only 35 percent of treatment group farmers. The KS test rejects the null hypothesis of distributional equality at the 5 percent level.

Although we find evidence of changes in timing of sales resulting from our information intervention, we only find an effect in the second year of the study. Thus, changes in timing cannot explain the positive treatment effect in the first year of the study.

B.4 Production decisions

Finally, we evaluate whether the price alerts had noticeable effects on farmers' production decisions. Given the nature of our study design—we provided farmers only with information on crops that they already grow—it is not immediately clear how or why the price alerts might affect production decisions, except perhaps by incentivizing farmers to produce more of a particular crop that they already grow.

To address this topic, we rely on the information collected in the annual surveys. Each annual survey gathered information on the acres of land devoted to cultivating each of the farmer's main two crops. We use this data to calculate the change in acres of land cultivated for a given crop, from baseline to follow-up, and from baseline to endline. We present results from two specifications: one with strata fixed effects as controls, and another with strata fixed effects along with individual and community controls. Results are presented in Table A4. In general, there appears to have been very little, if any, impact of the price alerts on acres of land dedicated to these crops. Interestingly, despite receiving higher prices for yam, treatment farmers don't seem to dedicate more land to the production of yam post-intervention. It is not clear whether this is because of land constraints, because the increase in prices was not enough to increase investments in land clearing and cultivation, or for some other reason.

Our endline survey also directly asked farmers whether they had made changes to their agricultural production over the previous season: whether they had started to grow a new crop, or whether they had started to grow more of any particular crop. We present the results of these questions in Table A5. One striking feature of the table is the fact that many farmers report changes to the crops they grow, and/or the proportion of land allocated to different crops. More than one-quarter of the sample reports starting to grow a new crop, and over 40 percent report growing more of an existing crop. However, for the most part, there are no significant differences between treatment and control. Taken together with Table A4, the results point strongly to the conclusion that the alerts did not lead to changes in farmers' production decisions during the period we studied.

Table A1: Impact of Esoko on bargaining for yams, follow-up survey

Dependent variable:	Farmer's offer price	Final price	Price expectation
Treatment	44.238***	24.181*	48.710***
	(16.385)	(12.842)	(14.238)
Quantity	9.343***	8.097***	10.878***
	(2.340)	(1.824)	(2.394)
Quantity^2	-0.088**	-0.078**	-0.151**
	(0.037)	(0.030)	(0.060)
Sold at home	-63.143*	-49.430**	-43.663
	(32.143)	(23.697)	(33.460)
Sold at local market	-4.054	-7.650	4.349
	(22.004)	(17.885)	(19.933)
Sold at urban market	69.395**	48.399*	49.376*
	(31.447)	(25.052)	(27.573)
N	497	499	424
R^2	0.324	0.320	0.325
Control group mean	333.28	251.53	289.53
Control group SD	131.54	108.44	113.31

Only cases where farmers bargained and made the first price offer are included (519 out of 546 observations). An additional 21 observations were omitted from the analysis due to missing information on quantity sold or place of sale. Prices are per 100 tubers and are in nominal Ghana cedis. Quantities are in hundreds of tubers. Regressions also include strata fixed effects and month of sale dummies. For place of sale, omitted category is farm gate. Huber-White robust standard errors clustered by community cluster are in parentheses.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

Table A2: Estimated effect of Esoko price alerts on where sold: Yam, Annual data

	Baseline		Follow-up		Endline	
	(1)	(2)	(1)	(2)	(1)	(2)
Panel A: Urban market sales						
Treatment	-0.009	-0.014	0.113***	0.075**	0.061	0.026
	(0.062)	(0.064)	(0.041)	(0.038)	(0.046)	(0.046)
R^2	0.106	0.171	0.246	0.278	0.128	0.187
Control group mean	0.23	0.23	0.17	0.17	0.17	0.17
Treatment group mean	0.25	0.25	0.31	0.30	0.26	0.25
Panel B: Local market sales						
Treatment	-0.083	-0.077	-0.079*	-0.071	-0.044	-0.031
	(0.070)	(0.065)	(0.046)	(0.048)	(0.048)	(0.046)
R^2	0.356	0.416	0.487	0.508	0.465	0.527
Control group mean	0.61	0.61	0.65	0.64	0.67	0.67
Treatment group mean	0.62	0.61	0.61	0.61	0.68	0.68
Panel C: Farm gate sales						
Treatment	0.019	-0.004	-0.004	-0.008	0.143**	0.109
	(0.078)	(0.078)	(0.064)	(0.064)	(0.071)	(0.072)
R^2	0.172	0.248	0.225	0.276	0.199	0.251
Control group mean	0.57	0.57	0.60	0.60	0.53	0.54
Treatment group mean	0.55	0.55	0.59	0.58	0.62	0.62
Strata FE	YES	YES	YES	YES	YES	YES
Pre-treatment outcome	N/A	N/A	YES	YES	YES	YES
Other controls	NO	YES	NO	YES	NO	YES
Observations	628	594	634	597	620	583

Dependent variable is binary indicator equal to 1 if the farmer reported selling in that location in the previous agricultural season. We use OLS to permit inclusion of strata fixed effects (a probit model was not possible since there was no variation in the dependent variable in some strata). The follow-up and endline regressions include a dummy for the pre-treatment outcome (sales in that particular location, from the baseline survey). "Other controls" include a set of individual farmer controls, and the community's distance to closest district market. Huber-White robust standard errors clustered by community cluster are in parentheses.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

Table A3: Effect of price alerts on yam prices: Monthly data through June 2012

	Price, level		Price, log	
	(1)	(2)	(1)	(2)
Treatment	7.143*	7.515*	0.059**	0.059**
	(3.829)	(3.844)	(0.028)	(0.026)
Sold in urban market		32.790***		0.247***
		(5.150)		(0.043)
Sold at farm gate		-8.670*		-0.046
		(4.902)		(0.033)
Sold at home		-6.194*		-0.053*
		(3.237)		(0.029)
N	2229	2224	2229	2224

Dependent variable is price per 100 tubers, reported in real August 2011 Ghana cedis. The omitted place of sale category is local markets. All regressions include month and strata fixed effects, individual and community controls. Robust standard errors are computed using 200 replication of block bootstrap.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

Figure A1: Impact of price alerts on timing of sales - CDFs of last month of sales

(a) Yam, 2012 season

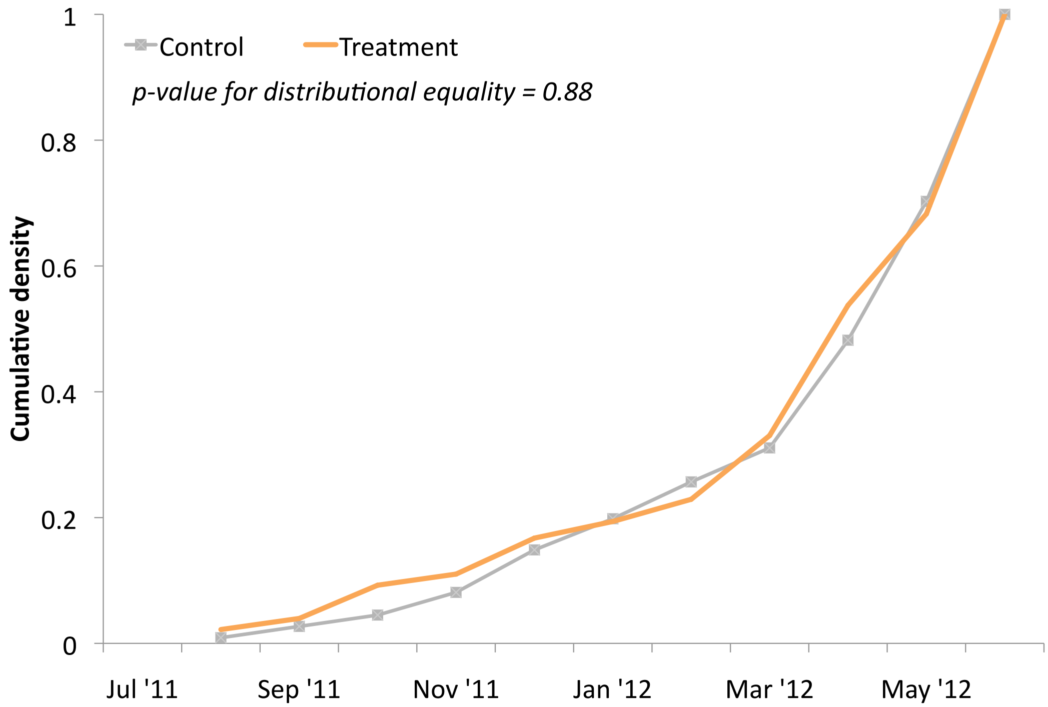

(b) Yam, 2013 season

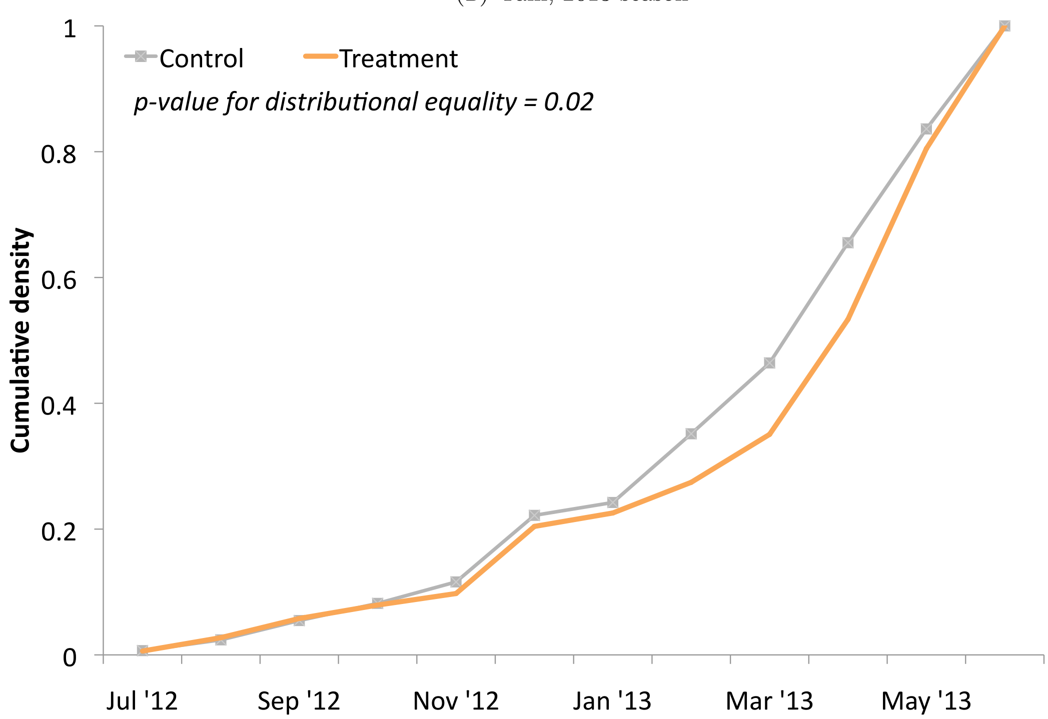

Notes: Figures show empirical CDF of month in which farmer sold the last of his crop for that season.

Table A4: Effect of alerts on change in land cultivated (acres): All crops, annual data

	Follow-up		Endline	
	(1)	(2)	(1)	(2)
Panel A: Yam				
Treatment	0.188	0.214	0.061	-0.007
	(0.273)	(0.252)	(0.367)	(0.347)
N	588	555	567	534
R^2	0.297	0.354	0.207	0.256
Panel B: Maize				
Treatment	-0.062	-0.091	0.241	0.180
	(0.209)	(0.206)	(0.321)	(0.306)
N	325	306	272	259
R^2	0.269	0.390	0.334	0.414
Panel C: Cassava				
Treatment	0.404	0.445*	0.626	0.594
	(0.258)	(0.227)	(0.414)	(0.403)
N	353	325	344	319
R^2	0.352	0.415	0.334	0.360
Panel D: Groundnut				
Treatment	-0.204	-0.085	-0.140	-0.130
	(0.228)	(0.266)	(0.196)	(0.258)
N	167	150	156	138
R^2	0.119	0.226	0.151	0.222
Strata FE	YES	YES	YES	YES
Other controls	NO	YES	NO	YES

Dependent variable is the change in the acres of land cultivated for a particular crop, relative to baseline. All regressions include strata fixed effects.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.

Table A5: Farmers self-reported changes in production, Endline survey

	N	C	T	T - C
Panel A: Growing a new crop				
Any crop	938	25.4%	29.1%	3.7%
Yam	311	3.6%	11.0%	7.4%
Maize	512	12.6%	10.7%	-1.8%
Cassava	517	3.0%	7.9%	4.8%
Groundnut	694	8.3%	7.0%	-1.3%
Rice	849	5.4%	7.5%	2.0%
Panel B: Growing more of any crop(s)				
Any crop	938	43.9%	45.0%	1.0%
Yam	613	40.2%	28.6%	-11.6%
Maize	382	5.3%	5.2%	-0.1%
Cassava	395	20.2%	14.8%	-5.5%
Groundnut	221	22.0%	10.8%	-11.2%*
Rice	76	7.7%	21.6%	13.9%

The survey asked farmers whether they had started to grow any new crops (or more of any existing crops) in the previous growing season (July 2012-June 2013). The "any crop" rows include all farmers. The crop-specific rows in Panel A only include farmers that did not grow that crop in previous years, according to the baseline and follow-up surveys. The crop-specific rows in Panel B only include farmers that grew that crop in the previous year, according to the follow-up survey. Standard errors of the difference are clustered at the community level.
*** Significant at 1% level. ** Significant at 5% level. * Significant at 10% level.