

Private Input Suppliers as Information Agents for Technology Adoption in Agriculture[†]

By MANZOOR H. DAR, ALAIN DE JANVRY, KYLE EMERICK, ELISABETH SADOULET,
AND ELEANOR WISEMAN*

Information frictions limit the adoption of new agricultural technologies in developing countries. Efforts to improve learning involve spreading information from government agents to farmers. We show that when compared to this government approach, informing private input suppliers in India about a new seed variety increases farmer-level adoption by over 50 percent. Suppliers become more proactive in informing potential customers and carrying the new variety. They induce increased adoption by those with higher returns from the technology. Being motivated by expanded sales offers the most likely motive for these results. (JEL D83, L14, L33, O13, O32, Q12, Q16)

Many people in poor countries rely on agriculture for their livelihoods. But they often use traditional technologies despite the existence of more productive alternatives. This puzzle has sparked extensive research to understand which barriers constrain technology adoption. Researchers have focused on credit and insurance market failures and on information frictions (Feder, Just, and Zilberman 1985; de Janvry, Sadoulet, and Suri 2017; Magruder 2018). Decades of research shows that learning plays an important role for adoption (Griliches 1957; Conley and Udry 2010; Fabregas, Kremer, and Schilbach 2019; Gupta, Ponticelli, and Tesei 2020; Cole and Fernando 2021).

Recognizing this, the public sector invests in agricultural extension—a process where government workers communicate information to selected contact farmers. This initial communication might trigger the flow of information through social networks. But the government system of outreach has not met expectations

*Dar: ICRISAT (email: manzoor.dar@icrisat.org); de Janvry: Department of Agricultural and Resource Economics, University of California, Berkeley (email: alain@berkeley.edu); Emerick: Department of Economics, Tufts University (email: kyle.emerick@tufts.edu); Sadoulet: Department of Agricultural and Resource Economics, University of California, Berkeley (email: esadoulet@berkeley.edu); Wiseman: King Center on Global Development and Center for Food Security and the Environment, Stanford University (email: ewiseman@stanford.edu). Benjamin Olken was coeditor for this article. This experiment on which this paper is based is registered in the AEA RCT registry as AEARCTR-0002121 (Dar et al. 2017). IRB approval was obtained from Tufts University (study number 1607005). The replication data and codes can be downloaded at the AEA Data and Code Repository (Dar et al. 2024). We acknowledge financial support from the J-PAL & CEGA's Agricultural Technology Adoption Initiative (ATAI). We are grateful for implementation support from the Odisha Department of Agriculture and Farmers' Empowerment. All errors are our own.

[†]Go to <https://doi.org/10.1257/app.20220037> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

(Farrington 1995; Anderson and Feder 2007). Outreach via the private sector may be one way to improve the flow of information to farmers, who rely on commercial input dealers for advice.¹ In India, farmers rely on input dealers for advice almost as often as they rely on their peers.² Motivated by profits, input dealers have incentives to provide information. As such, reaching farmers through the private sector could induce technology adoption.

We contrast two ways of disseminating information using a randomized experiment in Odisha, India. The main intervention provides information and seeds for testing to private input suppliers. We compare this approach with a control arm that replicates business-as-usual government extension. Many countries have piloted various forms of private agricultural extension (Rivera and Alex 2004). There is limited evidence on these reforms. We contribute by showing how information outreach can be more effective when delivered via the private sector compared to the government channel.

Involving private sector agents in providing public services can have ambiguous effects (Hart, Shleifer, and Vishny 1997). Input suppliers have motives to spread information about technologies they sell for profit. Repeated interactions with their clients may discipline them to provide high-quality recommendations.³ On the other hand, agro-dealers might provide lower quality when it is hard to verify (Bold et al. 2017; Ashour et al. 2019). Alternatively, they may recommend products that maximize their own profits instead of customer welfare, as researchers have found in other sectors (Hubbard 1998; Inderst and Ottaviani 2009; Mullainathan, Noeth, and Schoar 2012; Chen, Gertler, and Yang 2016; Anagol, Cole, and Sarkar 2017). Given these concerns, we seek to answer both whether informed agro-dealers increase adoption and whether they do so for the farmers that stand to gain the most from the innovation.

Our at-scale experiment took place in ten coastal districts of the Indian state of Odisha. Our sample consists of 72 blocks, covering an area with about 1.7 million farmers.⁴ We consider the dissemination of Swarna-Sub1, a new and profitable flood-tolerant rice variety.⁵ We partnered with the government extension service to support their conventional activities in 36 control blocks. This partnership included providing Swarna-Sub1 seed minikits to the contact farmers on whom they rely to spread information. It also involved carrying out large-scale “cluster” demonstrations where many farmers grow the new variety on contiguous plots of land and organizing farmer field days to share results from demonstrations. These are all

¹These dealers are in turn informed by the companies whose products they sell (Fites 1996) and by public agricultural agencies and research institutions (Wolf, Just, and Zilberman 2001).

²The 2018–2019 National Sample Survey finds that 20 percent of agricultural households rely on input dealers for technical advice (NSSO 2021). This is only slightly lower than the 23 percent that rely on other farmers and over 6 times the 3.1 percent that rely on extension agents.

³A potential downside of this approach is that markets for agricultural inputs in developing countries can be sparse. Aggarwal et al. (2022) show that travel costs to input suppliers play an important role in technology adoption for African smallholders. If these costs are too high, then few farmers will have contact with dealers, and solving information frictions on the supply side of technology could be less effective.

⁴Blocks are the relevant administrative units for agricultural extension. Blocks in the experiment have an average of 136 villages, and each block has its own local agricultural extension office.

⁵Previous work shows that this innovation is profitable for farmers. By reducing risk, it induces them to invest more in early-season inputs. Notably, it has no yield penalty in normal years (Emerick et al. 2016).

activities the government extension service would do with adequate resources, but we supported them as part of the experiment to make sure that they were carried out to the full extent and that our control group reflects business-as-usual activities at their best.

We provided the exact same quantity of seeds and the same information to input dealers in the 36 treatment blocks. We did not support any conventional extension activities, as was done in the control blocks. These dealers are highly local, small-scale businesses that sell seeds and often other inputs such as agrochemicals. They were free to choose how to use the demonstration seeds. The key distinction between this treatment and the standard mode of agricultural extension (the control) is that information constraints are being relaxed on the supply side with private agro-dealers rather than on the demand side with farmers. The treatment tries to leverage the economic incentives created by the marketplace for private businesses in transmitting information to their clients. Dealers in our sample receive the same profit margin for Swarna-Sub1 as they do for other seeds. Thus, increasing quantities sold, both currently and in future seasons, is the motive for dealers to promote new seeds to farmers.⁶

Turning to results one year later, we find that the dealer-based treatment increases adoption of Swarna-Sub1—the variety being introduced—by over 56 percent, that is, from 6.3 to 9.8 percent of farmers. Furthermore, the average farmer in the treatment blocks cultivates 69 percent more land with the variety.

Consistent with these farm-level results, the treatment triggered a supply-side response on the seed market. By the 2018 season, two years after we introduced the seeds, dealers in treatment blocks were about 59 percent more likely to have them in stock. We find some evidence that informing agro-dealers causes a change in local seed production. Treatment blocks produced 40–50 percent more Swarna-Sub1 seeds 3 years after the intervention.

Are farmers in the dealer treatment induced to adopt those with higher expected returns? To consider this, we look at heterogeneity according to past flood exposure—an important determinant of returns.⁷ The technology provides higher yields when crops are flooded relative to other types of rice grown by farmers. It leaves yields unchanged when there is no flooding. We find that the dealer treatment only increases adoption for farmers that are the most exposed to flooding. The treatment more than doubled adoption for the highest-risk farmers. It left adoption unchanged for lower-risk farmers. This finding suggests that dealers may consider the benefits to farmers when making suggestions.

Are these results short lived? The effects of information interventions in agriculture sometimes die out over time (Casaburi et al. 2019). Our main results come from a farm-level survey one year after the intervention. But the data we assembled on dealer-level supply and seed production cover up to four years after the intervention. We find no evidence that effects on these outcomes have declined

⁶Dealers in our setting principally sell seeds that are produced by the state-run seed corporation that fixes both wholesale and retail prices equally for all seed varieties.

⁷Flood tolerance is the key attribute of Swarna-Sub1. The technology provides higher yields when crops are flooded relative to other types of rice grown by farmers (Xu et al. 2006). It leaves yields unchanged when there is no flooding.

over time. The persistent effects are consistent with the dealer treatment leading to adoption by farmers with high expected returns. We would have expected rapid disadoption if dealers push the technology to farmers with low returns. Dealers increase profit by market size. This creates a powerful incentive to seek profit via long-term consumer satisfaction.

Turning to mechanisms, our treatment was designed to reduce a particular learning friction. Namely, high-powered incentives do not exist in the traditional model of agricultural extension. Agricultural extension agents in that model rely on nonincentivized communication between selected contact farmers and farmers in their social networks. By contrast, providing new information to profit-motivated agro-dealers results in their sharing this information with farmers.

The rest of our analysis shows supporting evidence for this mechanism where profit-motivated dealers give information to farmers. We show two separate pieces of evidence: one for information transmission and another for profit motives. Starting with information transmission, we sent “secret shoppers” to about 300 dealers. The purpose of these visits was to inquire about new rice varieties. This took place in the third season of the study—two years after the intervention. We find that the treatment changes what dealers say to potential customers. Dealers in treatment blocks are about 25 percent more likely to mention Swarna-Sub1 when listing the new varieties to consider. When asked for a specific recommendation, dealers in treatment blocks recommend older types of seeds at lower rates. In some cases, they are more likely to recommend trying Swarna-Sub1.

We then ran an experiment to test whether profit motives can explain why dealers communicate information that leads to greater adoption. In partnership with a local NGO, we revisited dealers in all blocks during the fourth season (2019). Each dealer was randomized into one of two treatments. In the first treatment, someone visited the dealer and asked which farmers, locations, and varieties would be best for a demonstration where farmers would cultivate a new variety and then the NGO would organize a meeting with other villagers to explain its attributes. Importantly, the name of the dealer giving the recommendation would be advertised during the meeting. In the second treatment, the same demonstration was described, but the NGO would not name the dealer and would collect the harvest after the demonstration and redistribute it as seeds to other farmers.

The first treatment presents a clear profit opportunity to the dealer. Making the dealer’s participation known can direct any demand created from the demonstration to them. The second treatment reduces these incentives by not broadcasting the dealer’s participation. It also redistributes seeds, which would reduce demand through the demonstration.

The treatment highlighting profit opportunities changes the advice given by dealers. It causes them to suggest different locations, types of farmers, and seed varieties. Starting with location, when presented with a candidate list of villages for the program, dealers in the profit motive treatment are more likely to suggest a village outside that list. They often suggest their own village. Dealers in this treatment also spend more time thinking of farmers to recommend. They are more likely to suggest neighbors or other people in their own village. Finally, the treatment causes dealers to recommend less common seed varieties.

Taken together, these findings are suggestive of a mechanism where our treatment first informs dealers. These dealers are then motivated by profit to pass this information along to farmers. But we cannot rule out other mechanisms. In particular, the treatment provided dealers with demonstration seeds. They could have passed those seeds to better connected farmers, who then shared the information. Additionally, dealers are usually farmers themselves. They may also be better connected, leading to more information transmission independent of their role as dealers.

We contribute by providing evidence on how private sector agro-dealers can improve the delivery of information to farmers. Most studies focus on improving the government system of spreading information. For instance, geographic information system (GIS) monitoring of government agricultural workers increases their effort (Dal Bó et al. 2021). Strategically selecting contact farmers and providing them financial incentives increases adoption (Beaman and Dillon 2018; BenYishay et al. 2020; BenYishay and Mobarak 2019; Beaman et al. 2021). Training the government-selected contact farmers has less of an impact (Kondylis, Mueller, and Zhu 2017). One common theme across these studies is the focus on government outreach to farmers through extension agents and contact farmers. We take a different approach by considering whether agricultural extension can leverage the private sector. Input suppliers have rarely been considered as information agents in agricultural markets.⁸ One related nonexperimental study in Niger finds that setting up demonstration plots with agro-dealers—as opposed to no demonstration plots at all—increases adoption of new seeds by farmers (Mamadou, Osei, and Osei-Akoto 2019).

More broadly, the private sector delivers public services in developing countries, including giving medical advice (Das et al. 2016; Kwan et al. 2022), distributing subsidized food (Banerjee et al. 2019), and providing services such as water and education (Galiani, Gertler, and Schargrotsky 2005; Romero, Sandefur, and Sandholtz 2020). In agriculture, full outsourcing of extension to the private sector has met challenges (Anderson and Feder 2004; Rivera and Alex 2004). The approach has proved much more complex to implement than expected, subject to collusion, and too expensive for farmers. Moreover, the fee-for-service approach only works when the advice is a market good—i.e., customized information that is excludable. Our intervention differs from full outsourcing. In our case, there is no contracting between public and private entities.⁹ Instead, we show how providing better information to private agro-dealers can outperform the provision of information through government channels.

The rest of this paper is organized as follows. Section I gives more information on the setting and outlines the experiment. Section II describes the data collection. Section III presents the main results on how assisting with the learning of input

⁸In the context of microcredit, Maitra et al. (2020) show that agricultural loans generate more benefits when private traders select recipients—rather than letting them self-select into group loans—perhaps due to incentive effects where private traders benefit themselves when loans cause farmers to harvest more output.

⁹Levin and Tadelis (2010) document how US municipalities are less likely to privatize services when performance measurement is difficult or there are holdup concerns. These issues are absent in our setting since we are not testing full outsourcing of agricultural extension.

dealers increases technology adoption by farmers, particularly those with the highest potential benefit. Section IV shows evidence that dealers spread information to their customers and that profits motivate them to do so. Section V concludes.

I. Background and Design of Main Experiment

This section starts by providing background information on the standard methods used in agricultural extension. It also gives a description of how the public sector delivers information to farmers in our particular study area. We then outline the design of our main experiment to compare these standard methods with the more business-oriented approach of using agro-dealers as information agents. The section concludes with a discussion of profits for agro-dealers and how they can only benefit from recommending a particular seed if it will increase their aggregate quantity sold, either currently or in the future.

A. Public Sector Agricultural Extension

Governments all over the world support agricultural extension services as a mode of information delivery. Ministries of agriculture typically have entire departments dedicated to providing these services. These departments oversee local offices that hire frontline extension agents whose role is to diffuse information about new agricultural technologies and practices to farmers. The specific techniques used by agents vary across contexts, but the basic methods are largely consistent, especially in poor countries. Agents usually work with selected “contact farmers” who are keen on trying new approaches and are best able to transmit knowledge to others in their social networks. They also organize farmer field days with cluster demonstration plots, where new technologies are implemented by multiple farmers to boost the diffusion of information.

In the context of our experiment, agricultural extension workers use these standard techniques. Each of the ten districts in the sample is organized into blocks, where a block has an average of 136 villages. Each block has an agricultural office that is led by a block agricultural officer (BAO). The BAO employs assistant agricultural officers (AAOs) and village agricultural workers (VAWs) who work in the field with farmers.

B. Experiment on Dealer-Based Extension

Our sample consists of 72 blocks in ten flood-prone districts of Odisha.¹⁰ We selected these areas because the promoted technology—a flood-tolerant rice variety called Swarna-Sub1—is most suitable for flood-affected areas.¹¹ The blocks in the sample represent around 20 percent of the blocks in the state.

¹⁰The districts are Bhadrak, Balasore, Cuttack, Ganjam, Kendrapara, Khorda, Jagatsinghpur, Jajpur, Nayagarh, and Puri.

¹¹The technology has been shown to benefit farmers by reducing yield losses when flooding occurs and therefore downside risk in any year, thereby increasing investment (Dar et al. 2013; Emerick et al. 2016).

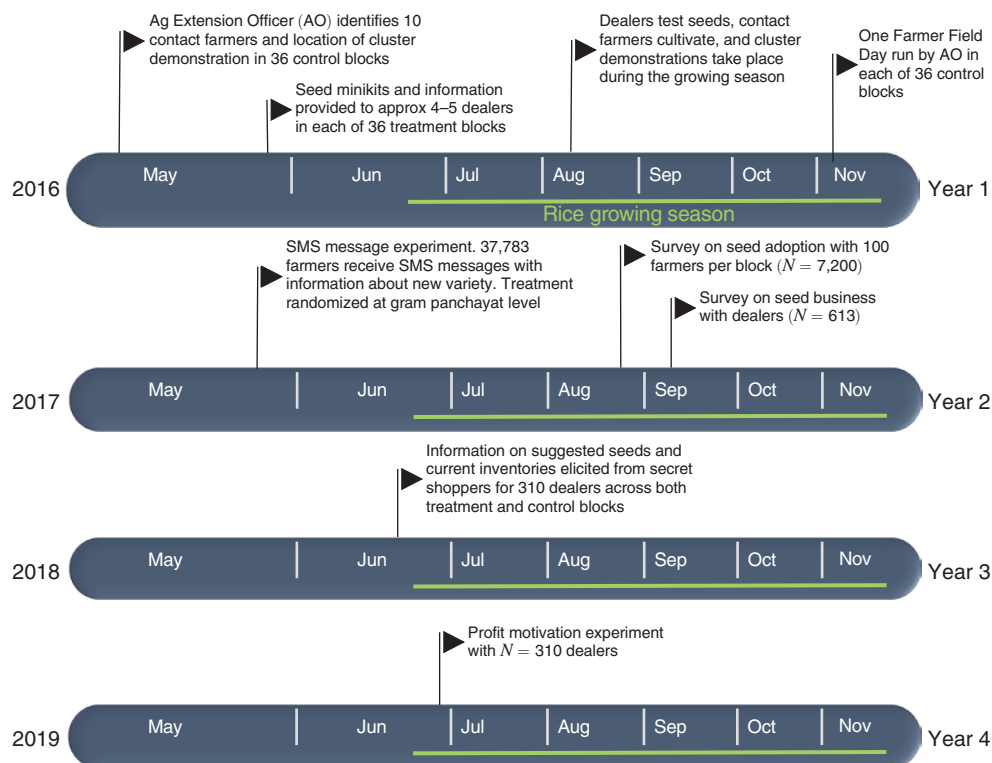


FIGURE 1. TIMELINE OF INTERVENTIONS AND DATA COLLECTION

We randomly assigned 36 of these blocks to the treatment group where agro-dealers were targeted to receive seeds and information. This randomization was stratified by district. The remaining 36 blocks serve as a comparison group where we supported the government extension service to carry out normal extension activities.

Figure 1 displays the timeline of these interventions. Starting in May 2016—about six to eight weeks before planting time—we partnered with the government’s extension service to introduce Swarna-Sub1 into control blocks. We did this in a way that mirrors three common practices in agricultural extension. First, field staff provided ten seed minikits of five kilograms (kg) each to the BAO, who then helped identify contact farmers to use the kits. The BAO chose two villages and five farmers in each village. Each kit contained only seeds for testing and some basic information about Swarna-Sub1. Our field staff then delivered the kits to the recommended farmers. Second, we provided another 150 kg of seeds to the BAO so that he could set up a cluster demonstration where the seeds would be used by several farmers on a contiguous set of plots. Based on seeding rates in the region, 150 kg allows for cultivation of 5–10 acres. The BAO chose where to do the demonstration and which farmers to target. Official government guidelines for organizing these clusters suggest that they be carried out in sites that are easily accessible for viewing by many farmers. Moreover, sites should be representative of average conditions in the area.

Third, we helped the BAO carry out a farmer field day in November—at the time right before harvest. The BAO selected the location of the field day and whom to invite. The purpose of the field day was for extension staff to train farmers about Swarna-Sub1 and share information from the demonstrations.

The objective of such an active control group is twofold. First, it ensures that each block is endowed with the same quantity of seeds. Therefore, the dealer-based treatment only differs in *who* received the new seeds and information. Second, the demonstrations and partnerships with contact farmers may not have taken place without our involvement. Forcing these activities to happen makes the treatment-control comparison more meaningful. Most importantly, it sets a higher bar for the dealer-based treatment by eliminating any possibility that the new technology would not be promoted by the government extension service.

Turning to the 36 treatment blocks, we obtained a list of 2,087 seed suppliers from the state Department of Agriculture. These include suppliers of two types: private seed dealers and Primary Agricultural Cooperative Societies (PACS). PACS are farmer groups that handle credit, seed supply, and procurement of output for farmers. We did not include them in the intervention, because their incentives are not the same as those of private dealers. Seed sales are usually handled by a member that is not the residual claimant on any profits from the sale. Despite being fewer in number relative to PACS, private dealers account for almost 60 percent of the seeds sold to farmers. The sample consists of 666 private dealers, 327 of which were located in the treatment blocks.

Armed with this list, our field staff entered each treatment block and located five dealers interested in receiving seed minikits and an informational pamphlet about Swarna-Sub1. In some blocks, fewer than five dealers were available. We provided additional seed to each dealer in these cases to guarantee that a full 200 kg (the same amount as control blocks) were introduced. The list provided by the Department of Agriculture did not have enough locatable dealers in some cases. In these circumstances, our field staff provided the seeds to other local agro-dealers.¹² Overall, seeds and information were provided to 151 dealers across the 36 treatment blocks.¹³ Of these, 119 were from the original list.

Once provided with seeds and information, the dealers were left alone to decide how to use them. We asked dealers about their intended uses. They overwhelmingly stated that they would use the seeds for testing on their own farms and would provide them to good customers for testing.¹⁴ Our intervention did not include any additional assistance to dealers. This approach differs from standard methods in agricultural extension where agents continually revisit their contact farmers. We allowed dealers to learn on their own because, in theory, they should be motivated to learn about a new product that could enhance their business. The goal of our

¹²The list of 666 dealers includes only those that are registered with the state, a prerequisite for selling seeds from the state seed corporation. The dealers not included in our list could have been in the process of renewing their license or only selling seeds produced by private companies.

¹³Two dealers in one control block were provided seeds by mistake. All of our analysis uses only the original random treatment assignment.

¹⁴Around 83 percent of dealers indicated that they would try some of the seeds on their own, while 63 percent indicated that some of the seeds would be provided to their good customers. Other less common responses were to provide them to family members (9 percent) and friends (24 percent).

treatment is to measure whether this motivation causes information to flow to farmers and ultimately increases adoption. Not intervening further ensures that our treatment effect is driven by any real-world incentives dealers have to learn rather than heavy monitoring by our partners.

We tracked these activities in both arms of the experiment. For control blocks, we took GPS coordinates of farmers chosen by the BAO for minikits. We collected the names and villages of the demonstration farmers from the BAO. Lastly, one of our team members attended 29 of the 36 field days and took GPS coordinates. In treatment blocks, we took photographs and GPS coordinates during seed delivery to dealers. Online Appendix Table A1 uses these GPS coordinates to show how proximity to the different extension sources varies by treatment. In particular, farmers in control blocks are 11.7 kilometers (km) from the nearest treated dealer. This distance falls to 4.36 km for treatment farmers.

Dealers in our sample are small business entrepreneurs. Some operate out of their homes, while others maintain shops in rural towns. Forty-four percent of dealers sell only seeds, with fertilizers and pesticides being the most common inputs sold by the other dealers. They are highly local. The average block in our study has 540 hectares (ha) of rice area per seed supplier.¹⁵ The median dealer in our data sells enough rice seed to cover roughly 162 ha, which implies that about 30 percent of area is planted with new seed each year. Farmers use seeds from their previous harvest for the remaining area. This creates an opportunity for dealers to expand sales by getting farmers to buy new seeds rather than use seeds from the prior year.

Turning to the second season (2017), we ran an SMS messaging experiment to compare our intervention with this “lighter touch” information treatment. The random delivery of SMS messages allows us to test whether our dealer treatment substitutes (or complements) basic knowledge that can be easily transmitted via ICT technology. Furthermore, it allows us to compare the direct effects of the two approaches.

The message informed farmers that Swarna-Sub1 is a new variety that is suitable for medium-low land in terms of elevation, matures in 145 days, and can tolerate up to two weeks of flooding. The message also stated that it was being produced by Odisha State Seed Corporation (OSSC) and could be available at local dealers. As a sampling frame, we obtained mobile numbers for 75,616 farmers that had registered for the state government’s Direct Benefit Transfer (DBT) scheme to obtain seed subsidies.¹⁶ These farmers are located across the 261 gram panchayats (an administrative unit usually consisting of around eight villages) that cover our main estimation sample, as outlined below. The SMS treatment was randomized at the gram panchayat level, resulting in messages being delivered to 37,783 of the names on the list.

¹⁵We arrived at this estimate by taking the total number of dealers and cooperatives in each block that are registered with the state seed corporation. Block-level rice area from 2016 to 2017 was available from the Odisha Directorate of Economics and Statistics.

¹⁶Beginning in 2016, the state government started providing seed subsidies in the form of payments back to farmers. Farmers were required to register, provide bank account details, and pay the full price at the time of seed purchase. The subsidy was then credited to their bank account after the transaction details had been entered into a mobile phone app by the seed dealer.

C. Motivation of Dealers to Recommend New Seeds

Dealers can be motivated to recommend a new seed either because it has a higher profit margin or because it will increase their quantity sold. The profit margin explanation cannot explain dealer behavior in our context. Around 84 percent of the seeds sold by dealers in our sample are produced by the state-run OSSC. As licensed agents of this company, dealers pay the same wholesale price for all rice varieties. State regulation fixes equal retail prices across varieties. Thus, the margin for dealers is identical across all types of varieties. This model, where state seed corporations play a major role in seed production and retailing, is common throughout India.¹⁷

As a result of these fixed markups, it is optimal for dealers to invest effort in recommending varieties that will increase their aggregate sales. Convincing farmers to shift from buying one variety to another is not profitable for the dealer, because of the equal margins. But getting farmers to purchase a new variety instead of using their own seeds of an older variety represents an increase in business. Put differently, dealers profit from recommending a new variety to promote seed replacement.

But dealers' incentives for providing advice go beyond encouraging seed replacement during a single season. Dealers sell seeds year after year. In this dynamic perspective, they benefit from making a recommendation today if it increases the size of their future business. This benefit can arise either because the same satisfied customers return or because farmers communicate to others that they learned about a profitable new variety from the dealer. Providing good advice today is a way for dealers to increase the quantities they sell during future seasons.

These features of the setting suggest that incentives are in place for dealers to promote improved seed varieties. The incentives operate through increasing current or future quantities, not through differential profit margins. Similar incentives do not exist for the public agricultural extension system. Government extension workers are not paid for performance, which may lower effort (Dal Bó et al. 2021). Even when government workers successfully influence contact farmers, these farmers might not spread information without incentives for themselves (Ben-Yishay and Mobarak 2019). Expanding dealers' knowledge about new seeds is thus meant to overcome the incentive problems that exist in traditional government extension.

II. Data Collection for Main Experiment

This section describes the experimental data for testing whether intervening with agro-dealers increases adoption by farmers. It also discusses the satellite data used to test whether adoption effects are larger for farmers with higher expected benefits from using the technology. We focus only on the data from the first two years of the study. We save the discussion of the additional data and the experiment on mechanisms for Section IV.

¹⁷For example, 80 percent of rice seeds in the state of Andhra Pradesh are purchased from retail outlets of the Andhra Pradesh State Seeds Development Corporation.

A. Survey on Farmer Technology Adoption

We anticipated that dealers and contact farmers would use the demonstration minikits for learning in 2016 and that any possible treatment effects could first be detected during year two (the 2017 season). Our main followup survey therefore took place in August–September 2017—around 15 months after the interventions. Its purpose was to measure adoption of seed varieties by rice farmers. To minimize measurement error, we timed the survey to occur right after planting.

Our sample consists of 7,200 farmers. These farmers were drawn from a random sample of 261 gram panchayats.¹⁸ Before drawing this sample, we excluded gram panchayats that had any village within 1.5 km of the block boundary.¹⁹ We removed these areas to reduce any interference caused by farmers possibly obtaining seeds from other blocks. The 261 sample gram panchayats had 75,616 farmers registered in the DBT program for seed subsidies. Using this database as a sampling frame, we randomly drew 100 farmers from each block (among the sample gram panchayats). These farmers are spread across 1,333 villages.²⁰ Online Appendix Figure A1 shows their geographic dispersion across the ten districts in the experiment.

Survey teams succeeded in locating and surveying 6,653 (92 percent) of the farmers. Of these, 93 percent were currently cultivating rice. Online Appendix Table A2 shows no significant differences in the probabilities of being surveyed or growing rice between treatment and control groups.

The survey focused on which seed varieties were currently being used for rice cultivation. Surveyors went through a list of 30 varieties and asked farmers which ones they were currently using and the amount of land being grown.²¹ In addition to these adoption data, we obtained information on contacts with agricultural extension agents during the last year, topics discussed during these conversations, whether the farmer had seen any seed demonstrations, and whether they had recently learned about Swarna-Sub1.

B. Data on Supply Responses

Any treatment effects on farmer-level uptake might occur simultaneously with supply responses by dealers.²² To measure this, we surveyed seed dealers around the same time as the farmer survey. We timed the survey to be in September so that seed purchases would be recently completed and easier to recall for dealers. Dealers were asked which varieties they carried for the 2017 season, how much of each was

¹⁸ We limited our data collection to a sample of gram panchayats to lower transportation costs for survey teams. The gram panchayats were identified using the 2011 population census of India (Office of the Registrar General and Census Commissioner India 2011).

¹⁹ Approximately 17.5 percent of the villages across the 72 blocks are within 1.5 km of another village in a different block.

²⁰ The farmer survey has almost no overlap with the dealers from treatment blocks. Using phone numbers of the treated dealers, we found only one of them in the farmer sample.

²¹ Swarna-Sub1—the variety introduced in the treatments—was twenty-fourth in this ordering. Asking about uptake in this way makes it less likely that responses reflect experimenter demands. Furthermore, farmers surveyed were not informed about the interventions that were carried out in their block a year earlier.

²² This need not be the case if there was already excess capacity of Swarna-Sub1 seeds or if farmers obtain seeds outside of formal markets, such as from friends or relatives.

sold, and whether they were selling seeds from private companies or from the state's seed corporation.

Our sample consists of 613 dealers from the list of dealers obtained prior to the experiment.²³ A large fraction could not be located or were no longer selling rice seeds. Specifically, 22.8 percent of them could not be reached. Of the 473 dealers located, 274 (58 percent) were selling rice seeds in the 2017 season. In results that follow, we show effects both for all dealers that were reached and for those that remained in the seed business. Online Appendix Table A3 shows that the likelihood of being located and the probability of selling rice seeds during the 2017 season are uncorrelated with treatment. Focusing on the treatment blocks, about 42 percent of the dealers surveyed received the intervention.

In addition to these dealer sales, we obtained data on the physical location of seed production. Seeds are grown by registered farmers that contract with the state to produce seeds that meet minimum certification standards. OSSC then collects, processes, and bags these seeds before selling them to farmers (via dealers and cooperatives) during the next season. The average block in our study had 32 seed growers per season from 2014 to 2019. We use records from a publicly available database that gives the location of each seed grower, the contracted area, the variety they produced, and the amount that was collected and processed (Odisha State Seed and Organic Products Certification Agency 2020).

Seed growers tend to be large farmers. They have incentives to produce the most profitable varieties for their land—just like other farmers.²⁴ As such, their production of a new variety depends on their being convinced of its potential. We therefore aggregate seed production at the block-season level and estimate the effect of the dealer treatment on the amount of Swarna-Sub1 produced in the block.

C. Flooding Exposure for Individual Farmers

Returning to farmer-level information, we use remote sensing data to approximate flooding risk. These data help us predict which farmers are expected to benefit the most from Swarna-Sub1. Being able to observe a key determinant of returns makes it possible to test for heterogeneous treatment effects according to a proxy for predicted benefits. More simply, is there a trade-off between intervening with private sector agents and a technology reaching the right people? Or does involving input suppliers in the diffusion of information cause technology to diffuse to high-return individuals?

We have GPS coordinates of the houses for 83 percent of the farmers that we surveyed in 2017.²⁵ These coordinates are matched to daily images of flooded areas from June to October for the period 2011 to 2017. We consider a household as

²³There were 53 dealers on our list of 666 that had no contact details, and thus we did not attempt to locate them.

²⁴The contracts with OSSC are on an acreage basis. OSSC and the grower agree on the variety, and OSSC purchases the output at a predetermined price. The grower pays for all of the inputs.

²⁵The likelihood of missing GPS coordinates is uncorrelated with treatment. A regression of observing GPS coordinates on treatment has a coefficient estimate of -0.018 and a t -statistic of 0.4 .

exposed to flooding on a given day if its house is within 1 km of any flooded area.²⁶ We then aggregate the total number of days of flood exposure across the seven years as a measure of flooding risk—and hence as a proxy for the return to Swarna-Sub1.

The online Appendix shows three characteristics of this variable. First, it varies substantially across the sample (online Appendix Figure A3). About 30 percent of households were not exposed to flooding. In contrast, 10 percent of households had flooding for 40 days or more. Second, this variation is partly driven by geographic characteristics. In particular, online Appendix Figure A4 shows that flooding is more frequent in lower-elevation areas that are closer to rivers. These correlations provide verification that our measure at least partly reflects underlying determinants of flooding risk—not just recent flood shocks. Third, farmers exposed to more flooding tend to be smaller, be poorer, and belong to low-caste social groups (online Appendix Table A4).

D. Descriptive Statistics

Table 1 shows descriptive statistics and verifies randomization balance. Panel A shows block-level characteristics, derived mostly from the 2001 census. Most notably, the blocks have around 136 villages and an average population of 110,000. Beyond these, treatment and control blocks look similar in a number of other characteristics, including local Swarna-Sub1 seed production, caste distribution of the population, and elevation.

Panel B shows characteristics of the respondents from our 2017 survey. These characteristics were collected after the treatment but are time invariant. Observables are mostly balanced for this sample that we use to estimate our main regressions.

III. Results of Dealer Extension Experiment

This section presents the results of the agro-dealer experiment. After outlining the estimation strategy in Section IIIA, Section IIIB shows that using dealers as information agents increases adoption by farmers. This finding is robust to different ways of measuring adoption and to including a battery of control variables. Section IIIC tests whether this treatment effect varies by exposure to flooding risk—which is highly correlated with expected returns. We turn to effects on the supply side in Section IIID. Particularly, we show effects on both dealer-level seed inventories and block-level production of Swarna-Sub1 seeds.

²⁶ A different study in one of the same districts collected GPS coordinates of both houses and rice plots (Emerick and Dar 2021, 2020). These data show that rice plots are within 1 km of the household almost 90 percent of the time (online Appendix Figure A2). The images of flooding extent are processed from MODIS by the DFO Flood Observatory (floodobservatory.colorado.edu). Each image has a spatial resolution of 250 meters. A pixel is classified as flooded on days when a ratio of the bands detecting surface water to land exceeds a numeric threshold. Using the GIS coordinates of each household, we calculate the distance between the household and the nearest flooded pixel for each day during June–October for 2011–2017.

TABLE 1—SUMMARY STATISTICS AND COVARIATE BALANCE

	Means		<i>p</i> -value
	Control	Treatment	
<i>Panel A. Block characteristics (N = 72)</i>			
Number villages	136.0 (64.04)	135.4 (60.52)	0.878
Population	110,687.2 (38,991.2)	120,997.2 (49,184.3)	0.296
Annual Swarna-Sub1 seed production	298.9 (421.2)	191.9 (264.6)	0.289
Share scheduled caste	0.209 (0.0492)	0.214 (0.0522)	0.617
Share scheduled tribe	0.0462 (0.0767)	0.0286 (0.0441)	0.136
Elevation (meters)	23.51 (24.43)	19.72 (19.35)	0.175
Literacy rate	0.727 (0.0726)	0.737 (0.0475)	0.425
Share agricultural	0.636 (0.141)	0.651 (0.0931)	0.537
Child sex ratio, 0–6 years	1.072 (0.0229)	1.069 (0.0214)	0.746
<i>Panel B. Farmer characteristics (N = 6,653)</i>			
Age	49.42 (11.12)	49.83 (11.38)	0.401
Years education	7.930 (4.467)	7.818 (4.533)	0.303
Below poverty line card	0.471 (0.499)	0.510 (0.500)	0.245
Female farmer	0.0794 (0.270)	0.0809 (0.273)	0.714
Scheduled tribe	0.0160 (0.125)	0.00569 (0.0752)	0.0336
Scheduled caste	0.150 (0.357)	0.153 (0.360)	0.853

Notes: Panel A shows means and standard deviations of block-level characteristics from the 2001 census of India (with exceptions of elevation and annual Swarna-Sub1 seed production). Elevation is calculated from satellite data, and the annual Swarna-Sub1 seed production is the average annual amount of seed processed from registered growers in the block from 2014 to 2016. It is measured in quintals (1 quintal = 100 kg). The literacy rate is defined as the number of literate individuals divided by the population older than six years old. The share agricultural is defined as the number of people working in agriculture divided by the working population. The child sex ratio is the number of male children 0–6 years old divided by the number of female children 0–6 years old. Panel B shows means and standard deviations of characteristics from our household survey. The *p*-values in column 3 are for the treatment variable in regressions of each characteristic on treatment and district (strata) fixed effects. Panel A uses robust standard errors, while panel B clusters errors at the block level.

A. Estimation

Our main analysis consists of farmer- or dealer-level regressions of outcomes on the block-level treatment indicator:

$$(1) \quad y_{ibd} = \beta \times \text{Treatment}_{bd} + \alpha_d + \varepsilon_{ibd},$$

TABLE 2—TREATMENT EFFECTS ON TECHNOLOGY ADOPTION

	Adoption				Acres	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.0348 (0.0193) [0.0949]	0.0324 (0.0147) [0.0709]	0.0312 (0.0146) [0.0739]	0.0641 (0.0311) [0.0649]	0.0611 (0.0251) [0.0340]	0.0640 (0.0251) [0.0250]
Dependent variable control mean	0.0634	0.0631	0.0631	0.0932	0.0926	0.0926
R^2	0.0280	0.0626	0.0617	0.0265	0.0449	0.0437
District fixed effects	X	X	X	X	X	X
Controls		X			X	
LASSO picked controls			X			X
Observations	6,653	6,599	6,601	6,653	6,599	6,601

Notes: The table shows the main treatment effects on adoption and acreage. All regressions use the data from the follow-up survey with farmers in August/September of 2017. The dependent variables are whether the farmer was currently using Swarna-Sub1 (columns 1–3) and the acreage cultivated with Swarna-Sub1 (columns 4–6). Columns 2 and 5 include all the control variables in the balance table. Columns 3 and 6 only include the control variables that are selected by an adaptive LASSO (ridge) regression. The standard error (clustered at the block level) is reported in parentheses below each point estimate. p -values calculated by randomization inference are reported in brackets below each standard error.

where i indexes farmers (or dealers), b indexes blocks, and d indexes districts. We include district fixed effects in all specifications because the treatment was stratified by districts. We cluster all standard errors by block.

B. Effect on Technology Adoption

Informing private input dealers and providing them with seeds to test leads to greater adoption by farmers when compared to conventional extension approaches used by the public sector. Table 2 shows this result. Each column gives treatment effects with clustered standard errors (parentheses) and p -values calculated by randomization inference (brackets).²⁷ Starting with column 1, farmers in treatment blocks are 3.5 percentage points more likely to adopt Swarna-Sub1 a year after the treatment, compared to farmers in control blocks. Given an adoption rate of 6.3 percent in the control group, this implies the treatment leads to a 56 percent increase in uptake. Columns 2 and 3 add controls for pretreatment covariates. Column 2 includes all the controls from Table 1, while column 3 uses the post-double-selection procedure in Belloni, Chernozhukov, and Hansen (2014) to select controls that predict either the outcome or the treatment. Adding controls does not affect the results.

The treatment caused acreage cultivated to increase: farmers in treatment blocks planted an average of 0.06 more acres with Swarna-Sub1 compared to farmers in control blocks, a 69 percent increase (column 4). Columns 5 and 6 show that acreage results are unaffected when adding control variables. This adoption effect operates on both the extensive and intensive margins: in addition to increasing the

²⁷For randomization inference, we use the randomization- t values described in Young (2019). Resampling is clustered at the block level, and we use 1,000 replications.

rate of adoption, the treatment increased the cultivated area of adopters (online Appendix Table A5).

Online Appendix Table A6 shows that the level of contact with extension agents or with cluster demonstrations is very low, even with our reinforced extension service in control blocks, and that farmers in treatment blocks were no less likely to be in contact with extensions workers or to have observed a demonstration of Swarna-Sub1 compared to control farmers.²⁸ In other words, we do not find evidence of displacement at the expense of other traditional channels.

Following up on the idea of displacement, we look at whether the treatment displaced other new varieties, potentially lowering welfare if it caused a shift away from high-quality seeds. We find no such evidence. Online Appendix Table A7 shows that the treatment had a negative effect on adoption of only two seed varieties—both of which were released over three decades ago. It does not appear that the increase in adoption caused by agro-dealers corresponds to a shift away from newly released technologies.

Finally, we find no evidence that the SMS messages increased adoption (online Appendix Table A8). They also did not change the effectiveness of the dealer treatment. The adoption gains from the dealer treatment cannot be obtained with a “lighter touch” SMS messaging intervention, at least in our context.

C. Heterogeneity

The evidence on average adoption rates shows that helping private agro-dealers learn is more effective than conventional approaches used in the public sector. A concern may be that as private agents, dealers optimize behavior based on their own profits,—in contrast with government extension agents who might factor in equity and may be better at targeting farmers who have high expected returns to adoption. It is, however, not obvious whether profit-maximizing dealers will deliver inferior targeting. Profit maximization strategies and farmers benefiting from adoption could coincide and may lead to similar outcomes, especially if we consider the repeated interactions between dealers and farmers over time.

In our context, being exposed to frequent flooding gives an easy-to-observe measure of potential returns—given the flood tolerance property of the variety.²⁹ We show that treatment dealers were successful at targeting Swarna-Sub1 to farmers who could benefit the most from the new technology, that is, farmers who live in flood-prone areas.

Figure 2 separates the sample by the satellite-based measure of past flooding and shows that treatment effects only exist in approximately half the sample where there were at least three flood days from 2011 to 2017. Conversely, the dealer treatment had little or no effect on adoption in the bottom half of the sample.

²⁸Only 5.7 percent of farmers report contact with the agricultural extension worker during the last year. This number is in line with other studies that showed low levels of contact between extension workers and farmers.

²⁹Our analysis focuses on flooding risk as one determinant of returns because it is measurable in our data. However, we acknowledge that the actual benefits to a given farmer depend on a number of factors, some of which are harder to observe, such as risk aversion.

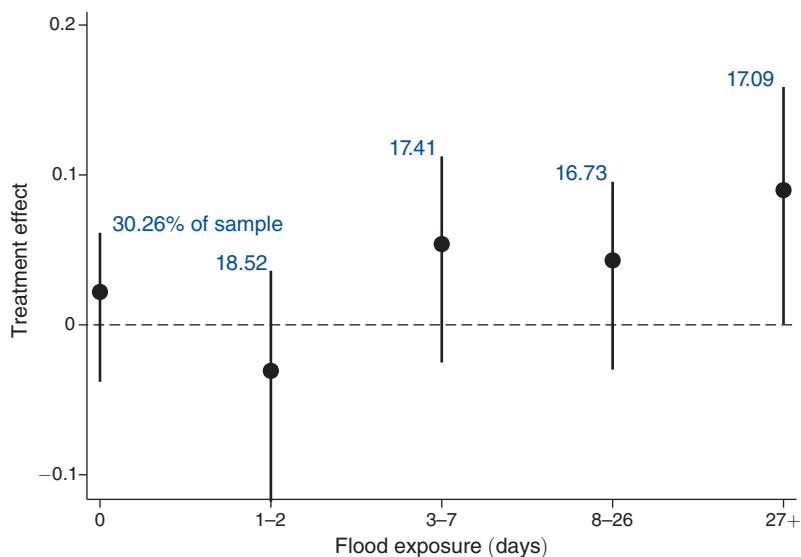


FIGURE 2. TREATMENT EFFECTS BY FLOOD EXPOSURE

Notes: The figure shows treatment effects from a single regression of adoption on separate treatment indicators for different levels of flood exposure and district fixed effects. The five bins of flood exposure correspond to households with no exposure from 2011 to 2017 and then an approximately equal division of households with at least one day of exposure. The percentage of observations in each bin is denoted in blue. The dots are the treatment effects of dealer-based extension, and the vertical lines denote 95 percent confidence intervals.

TABLE 3—HETEROGENEOUS EFFECTS BY FLOODING RISK

	Adoption		Acres	
	(1)	(2)	(3)	(4)
Treatment	0.00788 (0.0251)	-0.00404 (0.0218)	0.0260 (0.0426)	0.00194 (0.0389)
Treatment × Above-median risk	0.0563 (0.0330)	0.0589 (0.0304)	0.0909 (0.0590)	0.116 (0.0576)
Above-median risk	-0.0595 (0.0218)	-0.0644 (0.0242)	-0.0955 (0.0344)	-0.103 (0.0411)
District FE	Yes	Yes	Yes	Yes
Treat × Controls	No	Yes	No	Yes
Mean in control	0.063	0.063	0.089	0.088
p-value: Treatment + Treatment × High risk	0.018	0.014	0.009	0.002
Observations	5,536	5,495	5,536	5,495
R ²	0.036	0.072	0.033	0.054

Notes: The table shows heterogeneous treatment effects by flooding exposure. The dependent variable in columns 1–2 is an indicator for adopting Swarna-Sub1. Flooding risk is calculated by using satellite images from 2011 to 2017 (June–October) to count the total number of days when flooding was detected within 1 km of the farmer’s house. Above-median risk is a binary variable indicating a farmer that is above the median for the days of flood exposure. The dependent variable in columns 3 and 4 is the acreage cultivated with Swarna-Sub1. Columns 2 and 4 include interactions between all the control variables in Table 1 and the above-median risk variable. Standard errors are clustered at the block level.

TABLE 4—SEPARATE EFFECTS BY PAST FLOOD EXPOSURE

	High risk		Low risk	
	Adoption (1)	Acres (2)	Adoption (3)	Acres (4)
<i>Treatment</i>	0.0586 (0.0255) [0.0539]	0.111 (0.0437) [0.0170]	0.00222 (0.0256) [0.949]	0.0174 (0.0426) [0.743]
Dependent variable control mean	0.0357	0.0517	0.0831	0.116
R^2	0.0600	0.0416	0.0326	0.0320
District fixed effects	X	X	X	X
Observations	2,508	2,508	3,028	3,028

Notes: The table shows the main treatment effects on adoption and acreage separate for farmers with above- and below-median past flood exposure. All regressions use the data from the follow-up survey with farmers in August/September of 2017. The dependent variables are whether the farmer was currently using Swarna-Sub1 (columns 1 and 3) and the acreage cultivated with Swarna-Sub1 (columns 2 and 4). The standard error (clustered at the block level) is reported in parentheses below each point estimate. p -values calculated by randomization inference are reported in brackets below each standard error.

In Table 3, we show how the treatment effect depends on flooding risk. Two results stand out. First, control farmers from flood-prone areas are *less likely* to adopt Swarna-Sub1. The third row in the table shows that being a high-risk control farmer is associated with a 6 percentage point lower likelihood of adoption compared to low-risk control farmers. This estimate is merely a correlation. Farmers exposed to flooding differ in a number of ways that might directly influence adoption. Second, and more importantly, the dealer treatment was only effective in flood-prone areas—i.e., the interaction between treatment and flooding exposure is positive. Column 1 shows that the dealer treatment targets high-risk farmers increasing their adoption by 6.4 percentage points, while the effect of the treatment is only 0.8 percentage points for low-risk farmers (and not significant). The difference between the two treatment effects (the interaction term) is statistically significant at the 10 percent level. This interaction effect may be picking up other correlates of flood risk. Column 2 shows that the results are not sensitive to including interactions between the flooding risk variable and all the covariates in Table 1. We find similar results when looking at acreage in columns 3 and 4. The intervention only increased acreage of Swarna-Sub1 for farmers that were most exposed to flooding.

Table 4 shows analogous results where we split the sample according to flood risk. For farmers facing the most risk, the treatment increases adoption by 5.9 percentage points. This amounts to more than a doubling since only 3.6 percent of farmers adopted in the control group. Similarly, column 2 shows that the dealer intervention increases Swarna-Sub1 area by 0.11 acres for high-risk farmers. The intervention had no effect on adoption or acreage in places where flooding is less frequent (columns 3 and 4).

As another piece of evidence, online Appendix Table A9 shows that the average adopter in treatment blocks is more exposed to flooding. Specifically, they are more than twice as likely to be above the median in terms of flood exposure.³⁰

³⁰One possible concern is that we failed to obtain GPS coordinates for all farmers—they are missing for about 17 percent of the sample (1,110 respondents). In the online Appendix, we impute the locations of these houses

There is no evidence that informing dealers prioritizes adoption by the wealthiest farmers, which might have been expected if agro-dealers catered more to larger and wealthier farmers. In particular, online Appendix Table A10 shows that there is no treatment-effect heterogeneity according to farm size. Adoption is more likely by larger farmers, but this is equally true in treatment and control blocks. We also find no heterogeneity according to being below the poverty line or in a marginalized caste group.

D. Supply-Side Responses to the Treatment

Recall that we only treated a fraction of the dealers in each block. More precisely, 42 percent of sample dealers in treatment blocks received seeds and information (online Appendix Table A11). These dealers were not randomized. Hence, our dealer-level analysis compares all private dealers in treatment blocks to those in control blocks. We therefore capture any direct effect of receiving the seeds and information and any spillovers—which of course could be either negative or positive.

There is some evidence that the treatment caused dealers to increase the *availability* of Swarna-Sub1. Columns 1–4 in Table 5 show results from one year after the treatment (year 2). Focusing on all dealers—including those that were no longer operating—the treatment has a small positive effect on the likelihood of carrying Swarna-Sub1 at any time during the season (column 1) and the total amount the dealer reported selling throughout the year (column 2). But both of these estimates are very imprecise, partly due to some dealers no longer being in business. Among the subset of active dealers, those in treatment blocks were 6.2 percentage points more likely to carry Swarna-Sub1, a 17 percent increase (column 3). Column 4 shows that dealers in treatment blocks sold 3.7 additional quintals (1 quintal = 100 kg), which represents a 59 percent increase in volume sold. But again, while larger, neither of these results are close to statistically significant.

Anticipating on an intervention done in year 3 (and described below), we find large and precise effects on stocking behavior (column 5). Of dealers in control blocks, 19.3 percent had Swarna-Sub1 in stock when visited by the secret shopper.³¹ This increases by 11.4 percentage points (59 percent) in treatment blocks. This large effect is being observed two years after the treatment. It also comes from a direct observation of what the dealer had available on a certain day rather than an estimate from what they recalled after the season. Lastly, we collected secondary data from OSSC on which dealers carried Swarna-Sub1 during the 2020 season (year 5). Column 6 shows that the treatment effect persisted during that season.

This result could be driven by a number of things. First, it could come directly from the dealers that were treated and had their information sets updated. Online

using village locations in one of two ways. If we observe other households in that village, then we use the average latitude and longitude values from the observed households (603 farmers). If we observe no other households in the village, then we try to match the village to the 2011 census and use the village centroid as an approximate household location (323 farmers). Online Appendix Figure A5 shows that results are robust to including these observations in the flood-heterogeneity analysis.

³¹The probability of having Swarna-Sub1 available appears lower in year 3 for a couple of reasons: (i) availability was observed on a specific day when the shopper visited and not across the entire season, and (ii) visits by the secret shopper occurred later in the season when varieties were no longer in stock for some dealers.

TABLE 5—TREATMENT EFFECTS ON SUPPLYING SEEDS BY DEALERS

	Year 2: All		Year 2: In business		Year 3	Year 5
	Carry (1)	Quantity (2)	Carry (3)	Quantity (4)	In stock (5)	Carry (6)
Treatment	0.026 (0.052)	1.633 (1.676)	0.062 (0.083)	3.662 (2.656)	0.114 (0.057)	0.102 (0.060)
Dependent variable control mean	0.242	3.793	0.397	6.231	0.193	0.252
R^2	0.045	0.037	0.114	0.117	0.097	0.116
District fixed effects	X	X	X	X	X	X
Observations	473	472	274	273	280	252

Notes: The table shows treatment effects on Swarna-Sub1 inventories from a survey of dealers (columns 1–4), the secret shopper visit (column 5), and an online database with dealer-level inventories (column 6). Columns 1 and 2 are for the sample of dealers that were located and surveyed during year 2 (September 2017). Columns 3 and 4 are for the subset of those dealers that were actively in the seed business during that same season. Column 5 is for dealers that were visited in the secret shopper sample during year 3 (June 2018). Column 6 is for the 252 of these dealers for which we were able to obtain their license numbers. These numbers were matched to an online database with dealer inventories. The standard errors in each regression are clustered at the block level. The dependent variables are an indicator for whether the dealer reported carrying Swarna-Sub1 at any time during the season (columns 1 and 3), the total quantity sold throughout the season (columns 2 and 4), an indicator for whether the dealer had Swarna-Sub1 in stock when visited by the secret shopper (column 5), and an indicator for whether the dealer showed positive inventory in the online database at any time during the season (column 6).

Appendix Table A12 shows some evidence of this: there is a positive correlation between receiving the demonstration seeds and selling them during future years. Second, dealers talk to farmers. Any increase in the knowledge of farmers could spread to other dealers, not only those that were treated. Third, dealers were provided with several minikits for testing. They could have shared those in a way that increased local knowledge. We cannot distinguish between these effects in the analysis.

We next test whether the treatment changed the extent of local seed *production*. Our data here amount to six observations per block: three from the period before our treatment could have triggered a production response (2014–2016) and three from the posttreatment period (2017–2019). We therefore estimate block-level regressions of the amount of Swarna-Sub1 seed produced on treatment and district and year fixed effects.

We find evidence that treatment blocks produced more Swarna-Sub1 seeds after the experiment. Columns 1 and 2 of Table 6 show regressions using the total amount of seed production (in quintals) and its log. Seed production is highly skewed (online Appendix Figure A6).³² As a result, the point estimate is more precise in the log regression of column 2. It shows that the treatment led to a 47.9 percent increase in the amount of seed production, conditional on some production taking place. Columns 3 and 4 show that results become more precise when conditioning on average annual production during the 2014 to 2016 period. In column 3, treatment increased production by 79 quintals, or about 38 percent. Again the result is more precise in the log regression, as Column 4 shows that treatment

³²No Swarna-Sub1 seed is produced for just over 40 percent of block-year observations. At the same time, large amounts of seed are produced in a small number of blocks.

TABLE 6—TREATMENT EFFECTS ON LOCAL SEED PRODUCTION

	Time period: 2017–2019					
	Amount (1)	log (2)	Amount (3)	log (4)	Amount (5)	log (6)
Treatment	14.014 (81.938)	0.479 (0.262)	79.494 (52.494)	0.568 (0.209)		
2014–2016 production			0.736 (0.122)	0.002 (0.000)	0.736 (0.123)	0.002 (0.000)
Treatment × Year = 2017					76.218 (68.242)	0.575 (0.283)
Treatment × Year = 2018					53.207 (96.331)	0.740 (0.250)
Treatment × Year = 2019					109.057 (67.453)	0.300 (0.364)
Dependent variable control mean	209.511		209.511		209.511	
<i>p</i> -value: 2017 = 2018					0.670	0.544
<i>p</i> -value: 2017 = 2019					0.753	0.486
<i>p</i> -value: 2018 = 2019					0.667	0.314
<i>R</i> ²	0.149	0.391	0.435	0.532	0.436	0.536
District fixed effects	X	X	X	X	X	X
Year fixed effects	X	X	X	X	X	X
Observations	216	124	216	124	216	124

Notes: The table shows the effects of the dealer treatment on block-level seed production of Swarna-Sub1. Publicly available data on producer-level production of certified Swarna-Sub1 seeds were matched to the blocks in the experiment. The unit of observation in each regression is the block-season, where years range from 2014 to 2019. All columns are for the 2017–2019 period, one to three years after the intervention. The dependent variables are the annual amount of Swarna-Sub1 seed processed from growers in the block (columns 1, 3, and 5) and its logged value (columns 2, 4, and 6). Columns 3–6 control for the average annual production from 2014 to 2016 (the pre-period outcome). The standard errors in each regression are clustered at the block level.

blocks produced an average of 57 percent more Swarna-Sub1 seeds during the 3 years after the treatment.

Columns 5 and 6 show separate effects for the three seasons after the experiment. None of the year-specific effects are statistically distinguishable. This provides some evidence that the treatment effects have not deteriorated over time. The online Appendix helps visualize the results by showing cumulative distribution functions of seed production by treatment (online Appendix Figure A6). They show a noticeable rightward shift in the distribution for the treatment blocks, particularly in the top quintile of the distribution.

These results should not necessarily be interpreted as local production increasing to meet the growing demand of farmers. In fact, seeds are often processed outside of the block where they are grown, and can be sold anywhere in the state after processing. It seems more likely that intervening with agro-dealers caused more people to know enough about Swarna-Sub1 to cultivate it. This group includes seed producers, who are often large landholders, and can rely on some of the same sources of information as smaller farmers.³³

³³ Online Appendix Table A13 shows that there is a positive correlation between farmer-level adoption (from our survey) and block-level seed production. This is evidence that seed producers select varieties that are best suited for local conditions—and hence selected by farmers.

IV. Mechanisms and Motivations behind Dealers' Role in Increasing Technology Adoption

Our analysis up to this point shows that the information treatment targeted at agro-dealers causes more farmers to adopt new technology. That is, the approach of informing private agents on the supply side of technology can outperform standard approaches where frontline government workers interact directly with selected farmers. This section explores the mechanisms that are at play. More specifically, we first want to understand *how* dealers increase adoption. Section IVA shows evidence of one channel: dealers communicate information and recommendations to farmers. Secondly, we look into *why* treated dealers exhibit this behavior. We test whether profit opportunities motivate dealers when making recommendations (Section IVB).

A. Dealer Communication with Farmers

There are different ways dealers could have increased farm-level adoption. For one, they may advise clients to purchase Swarna-Sub1, playing an active and directed information-sharing role that goes beyond the traditional information sharing approach practiced by extension agents. Alternatively, they may have played more of an indirect role by informing people who are well connected, or even by giving demonstration seeds to better connected farmers.

We used the third year of the experiment (2018 season) to test whether dealers actively advised farmers about the new technology. Determining whether advice was given is not easy. Simply asking dealers whether they informed farmers and/or recommended Swarna-Sub1 likely suffers from experimenter demand effects.

We alleviate this concern with a unique strategy to elicit advice using “secret shoppers.” First, an enumerator visited dealers in both treatment and control blocks during the time when farmers usually buy seeds. The enumerator was someone the dealer had not seen before and who did not identify himself as part of the research. Then the enumerator followed a specific script to obtain advice from the dealer. The enumerator mentioned that his father from a nearby village was planning to cultivate rice and was looking for information on possible varieties to grow. The enumerator asked the dealer which varieties to consider without mentioning the name of any particular one.³⁴ Dealers usually mentioned several varieties—which we describe as the dealers “listing” varieties. If the dealer did not make a specific recommendation the shopper asked him which one he would recommend. We refer to this outcome as a dealer recommendation. If the dealer asked about the type of land, he was told medium-low in terms of elevation and, hence, risk of flooding. We also asked which varieties the dealer currently had in stock.

Given the costs of these visits and the scattered nature of our sample, we focused on the dealers that we reached during the previous year who were selling rice seeds. The sample consists of 310 dealers, 15 of which were not from our list obtained at

³⁴We phrased the question in these general terms to avoid priming the dealers to think about any particular variety.

TABLE 7—TREATMENT EFFECTS ON PROVIDING ADVICE TO SECRET SHOPPERS

	Control mean (1)	Treatment effect (2)	Treatment effect w/shopper FE (3)
Listed Swarna	0.834	−0.0157 (0.0353)	0.0129 (0.0337)
Listed Swarna-Sub1	0.510	0.120 (0.0740)	0.130 (0.0634)
Recommended Swarna-Sub1	0.297	0.0425 (0.0517)	0.0701 (0.0484)
Recommended Swarna	0.421	−0.116 (0.0467)	−0.131 (0.0495)
Listed Swarna-Sub1 and no stock	0.352	−0.00892 (0.0576)	0.00120 (0.0476)
Listed Swarna-Sub1 and stocked	0.159	0.129 (0.0514)	0.129 (0.0550)
Recommended Swarna-Sub1 and no stock	0.179	−0.0166 (0.0376)	−0.00249 (0.0419)
Recommended Swarna-Sub1 and stocked	0.117	0.0591 (0.0395)	0.0726 (0.0406)

Notes: The table shows treatment effects on the type of information provided by dealers when they were visited by a secret shopper during year 3 (June 2018). Each row shows results from a separate regression ($N = 280$) of that outcome variable on the block-level treatment indicator and strata fixed effects. Column 1 shows the control mean, while the coefficient estimate on the treatment indicator and its standard error are presented in column 2. Column 3 shows results when also including a shopper fixed effect. Listed Swarna and Swarna-Sub1 (rows 1 and 2) are binary variables for whether the dealer included that variety when listing good varieties to try. Recommended Swarna and Swarna-Sub1 (rows 3 and 4) are binary variables for whether the dealer recommended that variety when asked to make a specific recommendation. Rows 5–8 show effects on listing/recommending the varieties and whether or not they were currently in stock with the dealer. The standard errors in each regression are clustered at the block level.

the beginning of the study and 15 of which we did not reach. The sample for analysis therefore includes 280 dealers.

To assess whether dealers actively advised farmers, we look at how they interact with the secret shopper who visited their shops. If dealers do play an active role in promoting Swarna-Sub1 to farmers, dealers located in treatment blocks should be more likely to list and/or recommend Swarna-Sub1 compared to control blocks. If, instead, dealers do not play this active role, we should not expect to see differential recommendations from dealers between control and treatment blocks.

Table 7 shows results, where the rows are for separate outcomes. As with the earlier dealer-level results, these results are “intention to treat” since not all dealers in treatment blocks were informed. While the shoppers had a clear script to follow, it was impossible for the conversation to follow the same path for every single dealer. We therefore show results with and without fixed effects for the different shoppers.

Most dealers list Swarna as a popular variety, and this is not different between treatment and control blocks. However, dealers in treatment blocks were 12–13 percentage points more likely to list Swarna-Sub1 as a possibility to consider, a 25 percent increase given that control dealers list Swarna-Sub1 51 percent of the time. Furthermore, treatment dealers were about 4–7 percentage points more likely to recommend Swarna-Sub1—a large albeit nonsignificant effect. Not surprisingly, the

TABLE 8—TREATMENT EFFECTS ON ADOPTION FROM DEALERS AND OTHER SOURCES

	Dealers (1)	Other sources (2)
Treatment	0.00587 (0.0110) [0.650]	0.0289 (0.0136) [0.0619]
Dependent variable control mean	0.0223	0.0410
R^2	0.0107	0.0391
District fixed effects	X	X
Observations	6,653	6,653

Notes: The table shows the main treatment effects on adoption separately for whether the farmer reported getting the seeds from a dealer (column 1) or other sources (column 2). All regressions use the data from the follow-up survey with farmers in August/September of 2017. The dependent variable in column 1 is a binary variable equal to 1 for farmers that adopted Swarna-Sub1 and reported that they got seeds from dealers. The dependent variable in column 2 is a binary variable equal to 1 for farmers that adopted Swarna-Sub1 and reported they got seeds from any source other than dealers. The standard error (clustered at the block level) is reported in parentheses below each point estimate. p -values calculated by randomization inference are reported in brackets below each standard error.

increase in Swarna-Sub1 recommendations comes at the expense of recommendations for Swarna—the variety that is otherwise similar but does not offer flood tolerance. Indeed, being located in the treatment block reduces the likelihood that dealers recommend Swarna by 13 percentage points, a 31 percent decrease in recommendations for this older variety.

We previously showed that dealers in treatment blocks were more likely to be carrying Swarna-Sub1 at the time of these visits. The last four rows of the table show that listing and recommending Swarna-Sub1 go hand in hand with stocking it. In other words, the treatment causes dealers to both suggest the new variety to farmers and carry it in their shops. This is evidence that treatment dealers play a direct role in the increase of Swarna-Sub1 adoption by mentioning it to farmers.

An alternative interpretation of our results is that the intervention worked only because it increased supply, making it easier for farmers to obtain the seed. This explanation differs from a mechanism where the treatment triggers more spread of information. Our data allow for one test. Farmers obtain seeds from multiple sources, including dealers, other farmers, and agricultural cooperative societies. Our follow-up survey asked farmers where they obtained seeds. Table 8 shows evidence that the dealer intervention led to greater adoption from sources other than dealers. This provides evidence that the effect is not being driven only by increased supply from the treated dealers. Rather, outreach via the private sector leads to more informed farmers and thus increases demand.

B. Profit Opportunities as Motivation for Spreading Information

We conducted an additional small experiment during the fourth year (2019) to learn more about what motivates dealers to expend effort in making recommendations. As was discussed above, dealers can be motivated by profits from increasing

quantities sold, even when their margins for Swarna-Sub1 are the same as for other varieties. Beyond these additional profits, dealers could be motivated by prosocial preferences for other farmers' well-being, or they could simply want to be known as somebody who gives good advice.

We test the role of profit incentives by partnering with a local NGO that visited the dealers in our sample and asked for recommendations for a new program. More specifically, the NGO informed dealers that they would organize a seed demonstration in a village (located in the dealer's block) where farmers would cultivate a new variety and villagers would be invited to learn about it. They asked dealers for recommendations on the type of variety, which village to choose, and who would be the best farmers within the chosen village to select for cultivation. Dealers were randomly assigned to receive one of two different programs from the NGO. In the first group, the NGO made it clear that the dealer would have a profit opportunity. Dealers in that group were informed that their name would be displayed and advertised during the demonstration. Broadcasting the dealer's participation was chosen to increase the opportunity for him to profit from increased sales as a result of the demonstration. Dealers in the second group were asked for advice on a demonstration that was similar but had two key differences. First, the dealer's name would not be displayed as part of the demonstration. Second, the NGO explained to dealers that the harvest would be collected after the demonstration and redistributed as seeds to other farmers in the village. Distributing seeds was meant to further limit profit opportunities since it would reduce the need of farmers to buy seeds after the demonstration. We recorded the responses for four types of outcomes during the interview: dealers' effort invested in providing recommendations to the NGO, the varieties they suggest using, the village they recommend for the program, and, lastly, which farmers they suggest partnering with to carry out the demonstration.

Table 9 shows the results, which are divided into five groups. First, we find that dealers in the profit motivation treatment spent a nonsignificant 10 percent more time making recommendations (row 1). The second row shows that this increase in effort appears mostly when, in the conversation, dealers are asked which specific farmers to rely upon for the demonstration: a 14 percent increase in time invested on an average of just over 2 minutes picking farmers to recommend. Second, the profit motivation treatment seemed to shift dealers away from suggesting popular seed varieties and toward selecting them based on land type. The fourth row shows that when asked why they suggested a particular variety, 16 percent of dealers report doing so because it is locally popular; this falls by 7.1 percentage points, or 44 percent, in the treatment group.³⁵ While the estimate in the next row is not significant, the treatment seems to cause dealers to recommend varieties based on their agronomic characteristics, their suitability to the land, and how long they take to grow (duration).

Third, we find some evidence that the treatment changed how dealers recommend villages. Dealers were presented with a list of three randomly selected nearby villages to choose from. They were also given an option of recommending a village not

³⁵ Popular varieties mostly correspond to older seeds that farmers have been growing for a long time.

TABLE 9—RESULTS ON PROFIT MOTIVATION AND DEALER RECOMMENDATIONS

	Control mean	Estimate
<i>Timing</i>		
Total time spent (minutes)	5.093	0.526 (0.406)
Time spent picking farmers	2.200	0.303 (0.166)
<i>Variety selection</i>		
Number varieties recommended	2.893	-0.269 (0.212)
Selected popular	0.157	-0.071 (0.039)
Selected for duration/land type	0.179	0.078 (0.048)
<i>Village selection</i>		
Chose outside of list	0.371	0.127 (0.056)
Picked own village	0.493	0.069 (0.059)
Picked own village if chose outside	0.635	0.135 (0.086)
<i>Farmer selection</i>		
Chose because client	0.807	-0.092 (0.047)
Chose because villager/neighbor	0.121	0.091 (0.041)
<i>Business perception</i>		
Felt treatment would affect business	0.671	0.086 (0.052)

Notes: The table shows the effects of the profit motivation treatment on suggestions to an NGO for carrying out on-farm demonstrations. The treatment was informing dealers that their name would be affiliated with the demonstration, while the control group was not provided this information and was told that the seeds from the demonstration would be distributed to other villagers. Column 1 shows the mean of the outcome in the control group, while column 2 shows the point estimate and standard error from a regression of the outcome on the treatment and district fixed effects. The unit of randomization is the dealer. Therefore, robust standard errors are in parentheses for column 2.

on that list. The treatment increased the likelihood that dealers took this option by 12.7 percentage points (a 34 percent effect). Most dealers taking this option picked their own village. Among dealers that wanted to recommend a village not on their randomly selected list, 63 percent of the control group picked their own village. This increased by 13.5 percentage points for the treatment group, although the difference is not quite statistically significant.

Fourth, dealers were asked to identify three farmers to grow the seeds as part of the demonstration. About 81 percent of dealers report that they selected existing clients. This falls by 9.2 percentage points with the treatment. Linking a dealer's name to the recommendation seems to instead cause them to suggest farmers that are villagers or close neighbors. The dealer's ability to observe the demonstration

plot when it is cultivated by a close neighbor might explain this result. Dealers may put more value on having this ability when their name is linked to the demonstration.

Finally, dealers at the end of the visit were asked if they felt the demonstration would affect their business. Around 67 percent of the control group reported it would. This suggests that many dealers thought the demonstration would increase seed demand—even though their names would not have been identified. The treatment increased the perception that the demonstration would affect business by 8.6 percentage points, a 13 percent effect. This effect is modest, but it aligns with the other results showing how this subtle treatment changed some of the advice given by dealers.

These findings help interpret the results of our main experiment. They provide suggestive evidence that the advice given by dealers is motivated at least partly by their concerns about how it will affect their future profits. While helping to understand the incentives at play, this additional experiment has limitations. First, it does not mimic an everyday conversation between dealers and customers as closely as our secret shoppers did the previous year. Framing the conversation around an on-farm demonstration made it possible to experiment with the salience of profit opportunities. But it has the downside of not being the same type of conversation that would happen between a farmer and dealer. Second, the sample size is small, as it was limited to the active dealers from our original list. As such, some of the estimates are imprecisely estimated.

V. Discussion and Conclusions

This article provides evidence on how intervening on the supply side of agricultural input markets can be more effective than public agricultural extension services when providing information to promote the adoption of new technologies. Intervening with private input suppliers, that is, agro-dealers, would be a substantial transformation of the standard methods currently used. Government workers most often try to spread knowledge via direct contacts with selected farmers, expecting them to, in turn, diffuse information in their social networks. Much of current research on information constraints tries to identify ways of optimizing this approach. Our paper instead provides an empirical test of a different approach where information is transmitted to private input suppliers.

We find that informing private agro-dealers about a new and profitable seed variety, and giving them small amounts of demonstration seeds to test, causes farm-level adoption to increase by over 50 percent compared to the business-as-usual approach where government workers focus on outreach with selected farmers. Using the private sector approach increases adoption most among farmers with higher expected benefits from the technology. This improvement in targeting suggests that there is an alignment of incentives between dealers and farmers: dealers benefit from inducing farmers to adopt the right technology. We also found that our treatment triggers a supply response on the seed market. It causes dealers to be more likely to keep the seed in stock, and it increases local production of the seed.

Further evidence shows that these effects can be at least partly explained by dealers' actively advising farmers. Dealers in treated locations are more likely

to mention the new seed variety when asked what to grow by a “secret shopper.” Unpacking the incentives of agro-dealers is difficult and something we cannot do perfectly with our experiments. But we find some evidence in our second experiment that profit motives create incentives for agro-dealers to give advice.

Our findings thus show potential for a different approach to agricultural extension in developing countries: delivering information on the supply rather than the demand side of technology. There might be drawbacks to this type of approach in some contexts. In particular, if the technology is not a product that agro-dealers can sell for a profit, then using them as information agents may not increase adoption. But it appears that when profit motives exist, as in the case of our experiment, making input suppliers better informed can improve the practice of agricultural extension.

REFERENCES

- Aggarwal, Shilpa, Brian Giera, Dahyeon Jeong, Jonathan Robinson, and Alan Spearot.** 2022. “Market Access, Trade Costs, and Technology Adoption: Evidence from Northern Tanzania.” *Review of Economics and Statistics*. https://doi.org/10.1162/rest_a_01263.
- Anagol, Santosh, Shawn Cole, and Shayak Sarkar.** 2017. “Understanding the Advice of Commissions-Motivated Agents: Evidence from the Indian Life Insurance Market.” *Review of Economics and Statistics* 99 (1): 1–15.
- Anderson, Jock R., and Gershon Feder.** 2004. “Agricultural Extension: Good Intentions and Hard Realities.” *World Bank Research Observer* 19 (1): 41–60.
- Anderson, Jock R., and Gershon Feder.** 2007. “Agricultural Extension.” In *Handbook of Agricultural Economics*, Vol. 3, edited by Robert Evenson and Prabhu Pingali, 2343–78. Amsterdam: Elsevier.
- Ashour, Maha, Daniel Orth Gilligan, Jessica Blumer Hoel, and Naureen Iqbal Karachiwalla.** 2019. “Do Beliefs about Herbicide Quality Correspond with Actual Quality in Local Markets? Evidence from Uganda.” *Journal of Development Studies* 55 (6): 1285–306.
- Banerjee, Abhijit, Rema Hanna, Jordan Kyle, Benjamin A. Olken, and Sudarno Sumarto.** 2019. “Private Outsourcing and Competition: Subsidized Food Distribution in Indonesia.” *Journal of Political Economy* 127 (1): 101–37.
- Beaman, Lori, and Andrew Dillon.** 2018. “Diffusion of Agricultural Information within Social Networks: Evidence on Gender Inequalities from Mali.” *Journal of Development Economics* 133: 147–61.
- Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and Ahmed Mushfiq Mobarak.** 2021. “Can Network Theory-Based Targeting Increase Technology Adoption?” *American Economic Review* 111 (6): 1918–43.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen.** 2014. “Inference on Treatment Effects after Selection among High-Dimensional Controls.” *Review of Economic Studies* 81 (2): 608–50.
- BenYishay, Ariel, and A. Mushfiq Mobarak.** 2019. “Social Learning and Incentives for Experimentation and Communication.” *Review of Economic Studies* 86 (3): 976–1009.
- BenYishay, Ariel, Maria Jones, Florence Kondylis, and Ahmed Mushfiq Mobarak.** 2020. “Gender Gaps in Technology Diffusion.” *Journal of Development Economics* 143: 102380.
- Bold, Tessa, Kayuki C. Kaizzi, Jakob Svensson, and David Yanagizawa-Drott.** 2017. “Lemon Technologies and Adoption: Measurement, Theory and Evidence from Agricultural Markets in Uganda.” *Quarterly Journal of Economics* 132 (3): 1055–100.
- Casaburi, Lorenzo, Michael Kremer, Sendhil Mullainathan, and Ravindra Ramrattan.** 2019. “Harnessing ICT to Increase Agricultural Production: Evidence from Kenya.” Unpublished.
- Chen, Brian K., Paul J. Gertler, and Chun-Yuh Yang.** 2016. “Physician Ownership of Complementary Medical Services.” *Journal of Public Economics* 144: 27–39.
- Cole, Shawn A., and A. Nilesh Fernando.** 2021. “‘Mobile’izing Agricultural Advice Technology Adoption, Diffusion, and Sustainability.” *Economic Journal* 131 (663): 192–219.
- Conley, Timothy G., and Christopher R. Udry.** 2010. “Learning about a New Technology: Pineapple in Ghana.” *American Economic Review* 100 (1): 35–69.
- Dal Bó, Ernesto, Frederico Finan, Nicholas Y. Li, and Laura Schechter.** 2021. “Information Technology and Government Decentralization: Experimental Evidence from Paraguay.” *Econometrica* 89 (2): 677–701.

- Dar, Manzoor H., Alain de Janvry, Kyle Emerick, David Raitzer, and Elisabeth Sadoulet.** 2013. "Flood-Tolerant Rice Reduces Yield Variability and Raises Expected Yield, Differentially Benefiting Socially Disadvantaged Groups." *Scientific Reports* 3: 3315.
- Dar, Manzoor H., Alain de Janvry, Kyle Emerick, Elisabeth Sadoulet, and Eleanor Wiseman.** 2017. "Agro-Dealers and Adoption of Flood-Tolerant Rice in Odisha." AEA RCT Registry. <https://doi.org/10.1257/rct.2121-1.0> (accessed June 15, 2020).
- Dar, Manzoor H., Alain de Janvry, Kyle Emerick, Elisabeth Sadoulet, and Eleanor Wiseman.** 2024. "Replication Data for: Private Input Suppliers as Information Agents for Technology Adoption in Agriculture." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/E183346V1>.
- Das, Jishnu, Alaka Holla, Aakash Mohpal, and Karthik Muralidharan.** 2016. "Quality and Accountability in Health Care Delivery: Audit-Study Evidence from Primary Care in India." *American Economic Review* 106 (12): 3765–99.
- de Janvry, Alain, Elisabeth Sadoulet, and Tavneet Suri.** 2017. "Field Experiments in Developing Country Agriculture." In *Handbook of Economic Field Experiments*, Vol. 2, edited by Abhijit Vinayak Banerjee and Esther Duflo, 427–66. Amsterdam: Elsevier.
- Emerick, Kyle, and Manzoor H. Dar.** 2020. "Replication Data for: Farmer Field Days and Demonstrator Selection for Increasing Technology Adoption." Harvard Dataverse. <https://doi.org/10.7910/DVN/KSPAQI>.
- Emerick, Kyle, and Manzoor H. Dar.** 2021. "Farmer Field Days and Demonstrator Selection for Increasing Technology Adoption." *Review of Economics and Statistics* 103 (4): 680–93.
- Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor H. Dar.** 2016. "Technological Innovations, Downside Risk, and the Modernization of Agriculture." *American Economic Review* 106 (6): 1537–61.
- Fabregas, Raissa, Michael Kremer, and Frank Schilbach.** 2019. "Realizing the Potential of Digital Development: The Case of Agricultural Advice." *Science* 366 (6471): eaay3038.
- Farrington, John.** 1995. "The Changing Public Role in Agricultural Extension." *Food Policy* 20 (6): 537–44.
- Feder, Gershon, Richard E. Just, and David Zilberman.** 1985. "Adoption of Agricultural Innovations in Developing Countries: A Survey." *Economic Development and Cultural Change* 33 (2): 255–98.
- Fites, Donald V.** 1996. "Make Your Dealers Your Partners." *Harvard Business Review* 74 (2): 84–95.
- Galiani, Sebastian, Paul Gertler, and Ernesto Schargrofsky.** 2005. "Water for Life: The Impact of the Privatization of Water Services on Child Mortality." *Journal of Political Economy* 113 (1): 83–120.
- GADM.** 2018. "GADM, Version 3.6." Database of Global Administrative Areas. <https://gadm.org/data.html> (accessed May 6, 2018).
- Griliches, Zvi.** 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25 (4): 501–22.
- Gupta, Apoorv, Jacopo Ponticelli, and Andrea Tesei.** 2020. "Information, Technology Adoption and Productivity: The Role of Mobile Phones in Agriculture." NBER Working Paper 27192.
- Hart, Oliver, Andrei Shleifer, and Robert W. Vishny.** 1997. "The Proper Scope of Government: Theory and an Application to Prisons." *Quarterly Journal of Economics* 112 (4): 1127–61.
- Hubbard, Thomas N.** 1998. "An Empirical Examination of Moral Hazard in the Vehicle Inspection Market." *RAND Journal of Economics* 29 (2): 406–26.
- Inderst, Roman, and Marco Ottaviani.** 2009. "Misselling through Agents." *American Economic Review* 99 (3): 883–908.
- Kondylis, Florence, Valerie Mueller, and Jessica Zhu.** 2017. "Seeing Is Believing? Evidence from an Extension Network Experiment." *Journal of Development Economics* 125: 1–20.
- Kwan, Ada, Claire E. Boone, Giorgia Sulis, and Paul J. Gertler.** 2022. "Do Private Providers Give Patients What They Demand, Even If It Is Inappropriate? A Randomised Study Using Unannounced Standardised Patients in Kenya." *BMJ Open* 12 (3): e058746.
- Levin, Jonathan, and Steven Tadelis.** 2010. "Contracting for Government Services: Theory and Evidence from US Cities." *Journal of Industrial Economics* 58 (3): 507–41.
- Magruder, Jeremy R.** 2018. "An Assessment of Experimental Evidence on Agricultural Technology Adoption in Developing Countries." *Annual Review of Resource Economics* 10: 299–316.
- Maitra, Pushkar, Sandip Mitra, Dilip Mookherjee, and Sujata Visaria.** 2020. "Decentralized Targeting of Agricultural Credit Programs: Private versus Political Intermediaries." NBER Working Paper 26730.
- Mamadou, Adam, Robert Darko Osei, and Isaac Osei-Akoto.** 2020. "Impact of Reinforcing Agro Dealer Networks on Agricultural Productivity in Niger." *Journal of Agriculture and Sustainability* 13: 4.

- Mullainathan, Sendhil, Markus Noeth, and Antoinette Schoar.** 2012. "The Market for Financial Advice: An Audit Study." NBER Working Paper 17929.
- NSSO.** 2021. "National Statistic Survey Report No. 587: Situation Assessment of Agricultural Households and Land and Livestock Holdings of Households in Rural India, 2019." New Delhi: Government of India.
- Odisha State Seed and Organic Products Certification Agency.** 2020. "Seed Grower List." OSSOPCA. https://ossopca.nic.in/Publicinformation/Dist_growerer_Report.aspx (accessed April 10, 2020).
- Office of the Registrar General and Census Commissioner India.** 2011. "Census of India." <https://censusindia.gov.in/nada/index.php/>.
- Rivera, William, and Gary Alex.** 2004. "Privatization of Extension Systems: Case Studies of International Initiatives, Agriculture and Rural Development, Volume 2." World Bank Group Agriculture and Rural Development Discussion Paper 9.
- Romero, Mauricio, Justin Sandefur, and Wayne Aaron Sandholtz.** 2020. "Outsourcing Education: Experimental Evidence from Liberia." *American Economic Review* 110 (2): 364–400.
- Wolf, Steven, David Just, and David Zilberman.** 2001. "Between Data and Decisions: The Organization of Agricultural Economic Information Systems." *Research Policy* 30 (1): 121–41.
- Xu, Kenong, Xia Xu, Takeshi Fukao, Patrick Canlas, Reycel Maghirang-Rodriguez, Sigrid Heuer, Abdelbagi M. Ismail, Julia Bailey-Serres, Pamela C. Ronald, and David J. Mackill.** 2006. "Sub1A Is an Ethylene-Response-Factor-Like Gene That Confers Submergence Tolerance to Rice." *Nature* 442 (7103): 705–708.
- Young, Alwyn.** 2019. "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results." *Quarterly Journal of Economics* 134 (2): 557–98.